

# Design of an intermediate representation for query languages

## Internship report

Romain Vernoux, supervised by Kim Nguyễn (VALS team, LRI)

July 15, 2016

### Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Notations and conventions</b>	<b>3</b>
<b>3</b>	<b>Architecture</b>	<b>3</b>
<b>4</b>	<b>QIR calculus</b>	<b>5</b>
4.1	QIR data model . . . . .	5
4.2	QIR operators . . . . .	5
4.3	QIR expressions . . . . .	6
<b>5</b>	<b>QIR evaluation</b>	<b>7</b>
5.1	QIR reduction rules . . . . .	8
5.2	A measure for “good” output plans . . . . .	8
5.3	An exhaustive reduction strategy . . . . .	10
5.4	A heuristic-based reduction strategy . . . . .	11
<b>6</b>	<b>Experiments</b>	<b>18</b>
6.1	Example of code factorization: analytics queries . . . . .	18
6.2	Example of fragment grouping: dynamic queries . . . . .	20
6.3	Example of incompatibility: caching . . . . .	21
<b>7</b>	<b>Related work</b>	<b>21</b>
<b>8</b>	<b>Conclusion and future work</b>	<b>23</b>
	<b>Appendices</b>	<b>24</b>
<b>A</b>	<b>References</b>	<b>24</b>
<b>B</b>	<b>Internship timeline</b>	<b>25</b>
<b>C</b>	<b>Formal definition of occurrences</b>	<b>25</b>
<b>D</b>	<b>Proofs</b>	<b>26</b>
<b>E</b>	<b>Estimation of the heuristic fuel value</b>	<b>30</b>
<b>F</b>	<b>Source code and other resources</b>	<b>30</b>

# 1 Introduction

This internship is part of a joint collaboration between Laurent Daynes at Oracle Labs, Giuseppe Castagna and Julien Lopez (also a MPRI student) at PPS and Kim Nguyễn and myself at LRI. Each of these three groups worked on different aspects of the project and synchronized during regular meetings. In this report, I will do my best to focus on my contribution and explain only the necessary of the other parts of the project. I will use “I” to emphasize my work and “we” in formal definitions and technical sections.

## The general context

Data oriented applications, usually written in a high-level, general-purpose programming language (such as Java, Ruby, or JavaScript) interact with databases through a very coarse interface. Informally, the text of a query is built on the application side (either via plain string concatenation or through an abstract notion of statement) and shipped to the database. After its evaluation, the results are then serialized and sent back to the “application-code” where they are translated in the application language datatypes. Such roundtrips represent a significant part of the application running time. Moreover, this programming model prevents one from using richer query constructs, such as user-defined functions (UDF). UDFs are functions defined by the application developer, most of the time in a language more expressive than the query language, and used in a query (e.g., in a filter condition of a SQL query). While some databases also possess a “server-side” language for this purpose (e.g., PL/SQL in Oracle Database and PL/pgSQL in PostgreSQL), its integration with the very-optimized query execution engine is still minimal, resulting in notoriously slow query evaluations. The alternative is to evaluate UDFs in the application runtime, leading to additional application-database roundtrips and even poorer performance[8][7]. These problems, often referred to as *language-integrated query* issues, are in no way specific to relational databases and also affect the so-called NoSQL databases, which now all provide some form of declarative query language.

In this setting, Oracle Labs is developing Truffle[14], a framework for representing dynamic languages programs and functions as abstract syntax tree (AST) nodes with particular evaluation rules, and Graal[14], a just-in-time (JIT) compiler written in Java, leveraging Oracle’s experience from the Java JIT compiler and achieving high performance in evaluating Truffle nodes. A proof of concept, in the form of a full-blown JavaScript runtime prototype written entirely in Truffle and open-source projects for Ruby, R and Python runtimes[14], is already available. In principle, this framework could enable to efficiently evaluate UDFs written in high-level dynamic languages directly inside a database embedding a Java Virtual Machine (such as Oracle DB[15], Cassandra[11] or Hive[16]), thus bridging the aforementioned impedance mismatch gap, considerably decreasing the cost of roundtrips between query and UDFs evaluation, and providing the right level of abstraction for a deep integration in the database query optimization process.

## The research problem

Language-integrated queries and impedance mismatch were already research topics in the 1980s[1], although they have received increased attention with the popularization of web applications. Various solutions have already been proposed in the past[4] and constant demand for this matter has for instance lead Microsoft to include LINQ[12]—a framework integrating Microsoft SQL Server querying primitives in C#, among other languages—in .NET 3.5, and to continually add new features in the subsequent releases.

In our opinion, the imperfections of the existing solutions and the new perspectives opened by the Truffle/Graal project justified a reinvestigation of this (still trending) topic. Furthermore, taking in consideration the wide adoption of NoSQL databases, special care should be taken to propose a solution that is not specific to relational databases, and to the extent of our knowledge

such work has not been done yet. Thus, the problematic of the QIR project is articulated around the following questions:

1. How to draw the line between application code that should be evaluated in the application runtime, in the database native query processor and in the application language runtime embedded in the database?
2. How to represent queries in a way that is agnostic of the application language (for reusability) and the target database (to support relational and non-relational ones)?
3. How to rewrite this query representation to present an efficient query to the database?
4. How to translate this query representation to the database native language?
5. How to evaluate application (Truffle) code inside the database?

Although I took part in discussions revolving around all these questions (cf. Section 3), my internship was centered on questions 2 and 3. Questions 1 and 4 were investigated at PPS and question 5 at Oracle Labs.

## My contribution

A review of existing language-integrated query frameworks (cf. Appendix F) and results from a previous internship[13] highlighted that existing database query languages (including SQL) share high-level querying primitives (e.g., filtering, joins, aggregation) that can be represented by operators, but differ widely regarding the semantics of their expression language.

In order to represent queries in an application language- and database-agnostic manner (question 2 above), I designed a small calculus, dubbed “QIR” for *Query Intermediate Representation*, expressive enough to capture querying capabilities offered by mainstream languages. QIR contains expressions, corresponding to a small extension of the pure lambda-calculus, and operators to represent usual querying primitives (Section 4).

In the effort to send efficient queries to the database (question 3 above), I abstracted the idea of “good” query representations in a measure on QIR terms. Then, I designed an evaluation strategy rewriting QIR query representations into “better” ones (Section 5).

## Arguments supporting its validity

As an abstraction layer between application languages and databases, QIR guarantees the robustness of the entire framework regarding changes in these two tiers. But to provide formal evidence of the relevance of my solutions, I wrote an optimality proof for my evaluation strategy on a particular subset of QIR terms, with respect to the aforementioned measure (Theorem 6).

Additionally, I implemented a prototype of the QIR evaluation strategy and ran experiments showing that (i) the measure captures well the idea of “good” queries, (ii) rewritings have a low cost while enabling considerable speedup opportunities for the database and (iii) the evaluation strategy outputs optimal results well outside the scope of the QIR subset captured by the optimality proof (Section 6).

## Summary and future work

The ability to evaluate snippets from the application code inside the database opened new perspectives on the language-intergrated query problems. My approach allows a complete separation between application language and target database considerations, and to this extent is an interesting contribution even outside the scope of our project. Future work includes testing this concept with more combinations of application languages and databases, and observing the impact of such an architecture on real-world web applications.

## 2 Notations and conventions

**Trees** In this document, we will use an infix parenthesis notation for trees. The tree leaves are represented by their name and a tree rooted at a node  $N$  with  $n$  children  $S_1, \dots, S_n$  is denoted by  $N(r_1, \dots, r_n)$  where  $r_i$  is the representation of the subtree rooted at node  $S_i$ . For instance, if a tree  $T$  is composed of a node  $N$  which has two children  $P$  and  $Q$  such that  $Q$  has one child  $R$ , we will write  $T$  as  $N(P, Q(R))$ .

**Contexts** Contexts are trees in which some subtrees have been removed and replaced by “holes”. A hole is denoted by  $\square$ . For instance, a context  $C$  can be obtained by removing the subtree rooted at  $Q$  in  $T$  and will be denoted by  $N(P, \square)$ . Context holes can be filled with arbitrary subtrees. If a context  $C$  contains one hole,  $C[S]$  corresponds to the tree  $C$  where the hole has been replaced by the tree  $S$ . For instance, if  $C = N(P, \square)$  then  $T = C[Q(R)]$ . This definition generalizes to contexts with  $n$  holes using the  $C[S_1, \dots, S_n]$  notation. Notice that in this case, the context holes are ordered using a prefix (depth-first, left-to-right) traversal of the tree, which corresponds to the left-to-right order while reading the tree notation we use in this document. For instance, if  $C = N(\square, Q(\square))$  then  $T = C[P, R]$ .

**List comprehensions** Similarly to sets and the  $\{x \mid P\}$  notation for set comprehensions, we use the  $[x \mid P]$  notation for list comprehensions. List comprehensions are order-preserving, that is, a list  $[f(x) \mid x \in L]$  respects the order of  $L$  if  $L$  is a list. In this document, we will use them to build the list of a node’s children in a tree. For instance,  $0([x + 1 \mid x \in [1, 2, 3]])$  will stand for the tree  $0(2, 3, 4)$ . A node with an empty children list will be considered to be a leaf.

**Proofs** For space reasons, some proofs have been moved to the Appendices. When no proof is given directly below a lemma or a theorem, it can be found in Appendix D.

## 3 Architecture

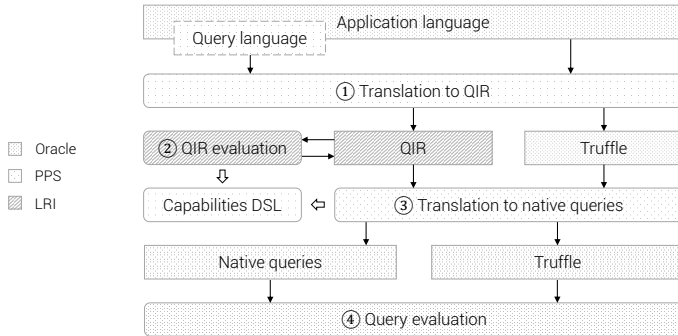


Figure 1: Architecture

In this Section, I present the global architecture of the project and quickly describe the purpose of the modules I did not work on.

Since the beginning of my internship, I took part in the design of the global architecture of the project, whose current state is illustrated in Figure 1. It works as follows.

1. Queries specified in the application language (possibly augmented with a special syntax for querying primitives) are detected in the application code. We refer to the result of this detection as the *query boundary* : code outside this boundary will be translated to Truffle nodes and evaluated in the application runtime, whereas code inside this boundary will be evaluated by the database. Inside the boundary, query parts that can be expressed in QIR are mapped to QIR constructs, and the remaining is propagated as Truffle nodes. This translation phase also maps the application language data model to the QIR data model. For now, we make the assumption that a same construct has identical semantics in the application language and in QIR (e.g., tuple navigations such as `user.name` return the same result in JavaScript and in QIR, and Ruby integers behave like QIR integers), and we allow a lossy translation when this is not the case. The QIR constructs will be described in Section 4. The translation phase, designed at PPS, was not the point of my internship

and will not be further described. To understand the remaining of this document, you can safely assume that query boundaries are specified by the application developer with a quote/antiquote syntax as in [3]. For example, in the following code snippet (written in an imaginary application language in which queries are specified using a SQL-like syntax),

```

1 f(id) := id = 1;
2 result := null;
3 <@ %result := select * from db("ads") as t where %f(t.id)) @>;
4 for t in result do ...

```

the `<@ @>` quotes delimit the query boundary, whereas the `%` antiquote annotate identifiers bound outside the query boundary. During the application execution, the `f` and `result` assignments are evaluated in the application runtime, the quoted code is translated to a QIR expression containing a Truffle reference `%f` and the dependencies of this quoted code (i.e., the Truffle definition of `f`) are passed, along with the QIR code, to the next module in order to be evaluated in the database.

2. QIR constructs are partially evaluated within the framework. The goal of this rewriting pass is to feed the database with queries that can be easily optimized by its query optimizer. This module relies on a formal description of the operators that the database can natively evaluate, in order to progressively “improve” the query representation with local reductions. Providing this description through a well-defined API is the first role of the “Capabilities DSL” (domain-specific language), thus maintaining the nice abstraction over the choice of target database. The partial evaluation of QIR expressions will be described in Section 5. This Capabilities DSL and description mechanism, designed at PPS, was not the point of my internship and will not be further described. To understand the remaining of this document, you can safely assume that it is an API providing information such as *This database supports filters, but not joins and a filter is compatible if the filter condition has the shape  $\lambda t. t.attr = x$  and if  $x$  is a constant value, but not if  $x$  is a free variable*.
3. After the rewriting pass, QIR expressions that can be expressed in the database native query language are mapped to native query constructs, and the remaining is translated to Truffle nodes. This translation phase also maps the QIR data model to the database data model. As in step 1, we make the assumption that a same construct has identical semantics in QIR and in the database query language, and we allow a lossy translation when this is not the case. Providing a way to translate compatible QIR operators into constructs of the database query language through a well-defined API is the second role of the Capabilities DSL, thus maintaining again the abstraction over the choice of target database. For the same reason as above, you can simply assume that this API provides a method mapping natively supported operators (according to the Capabilities DSL) to the corresponding query string in the database query language.
4. Native queries go through the database optimizer before their evaluation in the database query processor. Such queries may contain references to Truffle definitions for three reasons: (i) the query in the application language was referring to an identifier defined outside the query boundary, (ii) the quoted application code contained features that could not be translated into QIR, and (iii) the query representation in QIR contained features that could not be translated into the database query language. In all these cases, the Truffle runtime embedded in the database is called by the query processor to evaluate the code corresponding to these references. The result of the query is then returned to the host language. This backward translation maps database native data values to the QIR values (using again the Capabilities DSL translation API), then QIR values to the application language data model. Again, this evaluation mechanism, designed at Oracle Labs, was not the point of my internship and will not be further described in this document.

This architecture has the advantage of interleaving an abstraction layer, composed of QIR and Truffle, between the application layer and the database layer. Beyond added cleanliness,

this means that no work has to be done on existing database layer modules to support a new application language, and no work has to be done on existing application layer modules to support a new database.

## 4 QIR calculus

For the architecture described in Section 3 to work, QIR constructs have to be simple enough to be mapped to as many mainstream application language constructs and database language constructs as possible, but also expressive enough to represent interesting queries.

To design a QIR calculus satisfying these requirements, I first reviewed existing language-integrated query frameworks (cf. Appendix F), work that I had for the most part already done in a previous internship on a similar area[13]. I observed that existing database query languages all share high-level querying primitives (e.g., filtering, joins, aggregation), even though they support these features to various degrees. Moreover, the corresponding language constructs are almost always only a thin language layer hiding to various extents an algebra of operators inspired by the relational algebra. I also noticed the tension between database languages with a declarative descriptions of queries and application languages in which queries are built imperatively.

Therefore, in order to represent queries in a manner agnostic of the application and database languages, I designed QIR as a small extension of the pure lambda-calculus, made of (i) *expressions*, i.e., variables, lambdas, applications, constants, functions (conditionals, arithmetic operations, etc.) as well as list/tuple constructors and destructors, and (ii) *operators*, i.e., constructors representing the common querying primitives mentioned above. Thus, queries built imperatively in the application code can easily be translated to expressions manipulating operators whereas declarative queries correspond more to trees of operators, and one can go from the first representation to the second by partially evaluating the corresponding QIR term. The remaining of this section is a formal description of the QIR calculus.

### 4.1 QIR data model

QIR data objects are called *values*. The following grammar describes their syntax (for future use in this document) and hierarchy.

$$\begin{array}{ll}
 \langle \text{Value} \rangle & ::= \langle \text{Number} \rangle \\
 & \quad | \langle \text{String} \rangle \\
 & \quad | \langle \text{Bool} \rangle \\
 \langle \text{Number} \rangle & ::= \langle 64\text{-bit IEEE 754} \rangle \\
 \langle \text{String} \rangle & ::= \langle \text{UTF-16} \rangle \\
 \langle \text{Bool} \rangle & ::= \text{true} \\
 & \quad | \text{false}
 \end{array}$$

For now, the QIR data model only contains the necessary for a proof of concept. It will have to be extended with more primitive types for our framework to be used in real-world applications.

### 4.2 QIR operators

Operators represent computation on database tables. Each operator has *children* expressions (in parentheses) and *configuration* expressions (in subscript). As in the relational algebra, children provide the operator input tables whereas configurations describe the computation on table elements. The following grammar describes their syntax (for future use in this document) and hierarchy.

$$\begin{array}{ll}
 \langle \text{Operator} \rangle & ::= \text{Scan}_{\langle \text{Expr} \rangle}() \\
 & \quad | \text{Select}_{\langle \text{Expr} \rangle}(\langle \text{Expr} \rangle) \\
 & \quad | \text{Project}_{\langle \text{Expr} \rangle}(\langle \text{Expr} \rangle) \\
 & \quad | \text{Sort}_{\langle \text{Expr} \rangle}(\langle \text{Expr} \rangle) \\
 & \quad | \text{Limit}_{\langle \text{Expr} \rangle}(\langle \text{Expr} \rangle) \\
 & \quad | \text{Group}_{\langle \text{Expr} \rangle, \langle \text{Expr} \rangle}(\langle \text{Expr} \rangle) \\
 & \quad | \text{Join}_{\langle \text{Expr} \rangle}(\langle \text{Expr} \rangle, \langle \text{Expr} \rangle)
 \end{array}$$

These operators represent the following database computations:

- **Scan**<sub>table</sub>() outputs the list (in a non-deterministic order) of elements in the target database table denoted by the expression *table*.
- **Select**<sub>filter</sub>(*input*) outputs the list of elements *v* in the list corresponding to *input* such that *filter* *v* reduces to **true**.
- **Project**<sub>format</sub>(*input*) outputs the list of elements corresponding to *format* *v*, with *v* ranging in the list corresponding to *input*.
- **Sort**<sub>comp</sub>(*input*) outputs the list of elements *v* in the list corresponding to *input* ordered according to *comp* *v* ascending.
- **Limit**<sub>limit</sub>(*input*) outputs the *limit* first elements of the list corresponding to *input*.
- **Group**<sub>eq,agg</sub>(*input*) outputs the list of elements corresponding to *agg* *g* for each group *g* in the partition of the elements *v* in the list corresponding to *input*, according to *eq* *v*.
- **Join**<sub>filter</sub>(*input*<sub>1</sub>, *input*<sub>2</sub>) outputs the join of the elements *v*<sub>1</sub> in the list corresponding to *input*<sub>1</sub> and the elements *v*<sub>2</sub> in the list corresponding to *input*<sub>2</sub>, such that *filter* *v*<sub>1</sub> *v*<sub>2</sub> reduces to **true**.

For the same reason as above, QIR operators are still limited. Nevertheless, they already capture a consequent part of the relational algebra.

### 4.3 QIR expressions

QIR expressions are the core of the QIR calculus. The following grammar describes their syntax (for future use in this document) and hierarchy.

$\langle \text{Expr} \rangle$	$::=$	$\langle \text{Variable} \rangle$ $\mid$ $\langle \text{Lambda} \rangle$ $\mid$ $\langle \text{Application} \rangle$ $\mid$ $\langle \text{Constant} \rangle$ $\mid$ $\langle \text{ValueCons} \rangle$ $\mid$ $\langle \text{ValueDestr} \rangle$ $\mid$ $\langle \text{ValueFun} \rangle$ $\mid$ $\langle \text{BuiltinFun} \rangle$ $\mid$ $\langle \text{DataRef} \rangle$ $\mid$ $\langle \text{TruffleNode} \rangle$ $\mid$ $\langle \text{Operator} \rangle$	$\langle \text{TupleCons} \rangle$	$::=$	$\text{'tnil'}$ $\mid$ $\text{'tcons' } \langle \text{String} \rangle \langle \text{Expr} \rangle \langle \text{Expr} \rangle$
$\langle \text{Lambda} \rangle$	$::=$	$\text{'}\lambda\text{' } \langle \text{Variable} \rangle \text{'.' } \langle \text{Expr} \rangle$	$\langle \text{ValueDestr} \rangle$	$::=$	$\langle \text{ListDestr} \rangle$ $\mid$ $\langle \text{TupleDestr} \rangle$
$\langle \text{Application} \rangle$	$::=$	$\langle \text{Expr} \rangle \langle \text{Expr} \rangle$	$\langle \text{ListDestr} \rangle$	$::=$	$\text{'destr' } \langle \text{Expr} \rangle \langle \text{Expr} \rangle \langle \text{Expr} \rangle$
$\langle \text{Constant} \rangle$	$::=$	$\langle \text{Value} \rangle$	$\langle \text{TupleDestr} \rangle$	$::=$	$\text{'tdestr' } \langle \text{Expr} \rangle \langle \text{String} \rangle$
$\langle \text{ValueCons} \rangle$	$::=$	$\langle \text{ListCons} \rangle$ $\mid$ $\langle \text{TupleCons} \rangle$	$\langle \text{ValueFun} \rangle$	$::=$	$\text{'-'} \langle \text{Expr} \rangle$ $\mid$ $\langle \text{Expr} \rangle \text{'and' } \langle \text{Expr} \rangle$ $\mid$ $\langle \text{Expr} \rangle \text{'if' } \langle \text{Expr} \rangle \text{'then' } \langle \text{Expr} \rangle \text{'else' } \langle \text{Expr} \rangle$ $\mid$ ...
$\langle \text{ListCons} \rangle$	$::=$	$\text{'nil'}$ $\mid$ $\text{'cons' } \langle \text{Expr} \rangle \langle \text{Expr} \rangle$	$\langle \text{BuiltinFun} \rangle$	$::=$	$\text{'avg' } \langle \text{Expr} \rangle$ $\mid$ ...
			$\langle \text{DataRef} \rangle$	$::=$	$\text{'db.table\_name'}$
			$\langle \text{TruffleNode} \rangle$	$::=$	$\text{'truffle<id>'}$

Lambda-abstractions (resp. applications, constants) represent functions (resp. function applications, constants) of the application language. The constructor for lists takes an expression for the head and an expression for the tail. Tables are represented by lists, since the output of queries might be ordered. Tuples are constructed as a list of mappings: the constructor for tuples takes a string for the mapping name, an expression for the mapping value and an expression for the tail of the mapping list. The list destructor has three arguments: the list to destruct, the term to return when the list is **nil** and a function with two arguments  $\lambda h.\lambda t.M$  to treat the case when the list has a head and a tail. The tuple destructor has two arguments: the tuple to destruct and the attribute name of the value to return. Finally, built-in functions represent



common database functions (such as aggregation) for easier recognition and translation into native languages. The formal definition of expression reductions is given in Section 5.

In this document, we will sometimes use the syntactic sugar `let  $x = M$  in  $N$`  for `( $\lambda x. N$ )  $M$`  but the `let...in` constructor is not part of the expression language. Similarly, we will sometimes use the notation  `$t.attr$`  for `tdestr  $t$  "attr"`, `not  $x$`  for `if  $x$  then false else true` and `let rec  $x = M$  in  $N$`  as a shortcut for the application of a fixpoint combinator.

## 5 QIR evaluation

As discussed in the beginning of Section 4, the purpose of the QIR partial evaluation mechanism is to transform the QIR representation of a query written in the (possibly imperative) style of the application language into a QIR representation of the same query that is easier to translate to the database native language and offers more optimization opportunities to its optimizer. Today's databases use a tree of algebra operators (plan) to represent computation. Optimizations consist in commuting operators in the plan or applying local tree rewritings. The plan evaluation usually corresponds to (a variation of) a bottom-up evaluation of each operator, where a particular implementation is chosen for each operator for best performance. Therefore, for a same computation, such database engines benefit from working on a large, single plan instead of multiple small plans using the result of each other, since (i) more optimizations become available, (ii) more statistics can be gathered to pick the best implementation for each operator and (iii) no intermediate result needs to be materialized and transferred.

Consider the following example of query (written in an imaginary application language in which queries are specified using a SQL-like syntax).

```

1 result := null;
2 <@ users := db("users"); ads := db("ads");
3 filter(x,y) := (x.id = y.userid) and (y.timestamp > 1234);
4 prettify(s) := ...;
5 %result = select prettify(name) as u_name, descr as a_descr
6           from users as u, ads as a where filter(u,a) @>
7 for t in result do ...

```

The quoted code contains a query as well as the definition of two UDFs: `filter`, which is a quite simple filter function and `prettify`, which contains application languages features that cannot be translated to QIR. A direct translation of the quoted code to QIR would be as follows.

```

1 let users = db.users in let ads = db.ads in
2 let filter =  $\lambda x. \lambda y. (tdestr\ x\ "id") = (tdestr\ y\ "userid")\ and\ (tdestr\ y\ "timestamp" > 1234)$  in
3 let prettify = truffler<0> in
4 Project $\lambda t. tcons\ "u\_name"\ (prettify\ (tdestr\ t\ "name"))\ (tcons\ "a\_descr"\ (tdestr\ t\ "descr")\ tnil)\ ($ 
5   Join $\lambda u. \lambda a. filter\ u\ a\ (Scan_{users}(), Scan_{ads}())$ )

```

Notice that the definition of `prettify` had to be translated to Truffle nodes (not shown here) and is represented in the QIR code by a Truffle reference. Assuming that the target database can only natively evaluate operators (e.g., SQL), the direct translation of the above QIR code to the database language would lead to a query in which all operators issue calls to UDFs, greatly impacting performances. Moreover, with no information about `filter`, the database optimizer could not decide to push the selection `timestamp > 1234` below the join or choose an index-based implementation for the join. In fact, a better QIR representation for this query is as follows.

```

1 Project $\lambda t. tcons\ "u\_name"\ (truffler<0>\ (tdestr\ t\ "name"))\ (tcons\ "a\_descr"\ (tdestr\ t\ "descr")\ tnil)\ ($ 
2   Join $\lambda u. \lambda a. (tdestr\ u\ "id") = (tdestr\ a\ "userid")\ and\ (tdestr\ a\ "timestamp" > 1234)\ (Scan_{db.users}(), Scan_{db.ads}())$ )

```

This second version can be obtained from the first by a reduction of the QIR code. As we will see later in this section, applying classical reduction strategies (e.g., call-by-value, call-by-name, lazy evaluation) until a normal form is reached does not satisfy our requirements, since contracting



some redexes might scatter parts of a same query. Instead, we will characterize the shape of “good” QIR code with a measure on QIR expressions in Section 5.2, then investigate reduction strategies guided by this measure in Sections 5.3 and 5.4.

## 5.1 QIR reduction rules

In this section, I formally define the reduction of QIR expressions, then state two useful properties on the QIR calculus equipped with this reduction.

**Definition 1.** *The capture-avoiding variable substitution in QIR expressions (cf. Section 4.3) is defined as in the pure lambda-calculus. The substitution of a variable  $x$  in  $e_1$  by an expression  $e_2$  is denoted  $e_1\{e_2/x\}$ .*

**Definition 2.** *The reduction rules for QIR expressions consist in the  $\beta$ -reduction rule augmented with the  $\delta$ -reduction rules for destructors and  $\rho$ -reduction rules for primitive types.*

$(\lambda x. e_1) e_2$	$\rightarrow_\beta e_1\{e_2/x\}$
$\text{destr nil } e_{\text{nil}} e_{\text{cons}}$	$\rightarrow_\delta e_{\text{nil}}$
$\text{destr (cons } e_{\text{head}} e_{\text{tail}}) e_{\text{nil}} e_{\text{cons}}$	$\rightarrow_\delta e_{\text{cons}} e_{\text{head}} e_{\text{tail}}$
$\text{tdestr (tcons "name1" } e_{\text{val1}} e_{\text{tail}}) \text{"name1"}$	$\rightarrow_\delta e_{\text{val1}}$
$\text{tdestr (tcons "name1" } e_{\text{val1}} e_{\text{tail}}) \text{"name2"}$	$\rightarrow_\delta \text{tdestr } e_{\text{tail}} \text{"name2"}$
$\text{if true then } e_1 \text{ else } e_2$	$\rightarrow_\rho e_1$
$\text{if false then } e_1 \text{ else } e_2$	$\rightarrow_\rho e_2$
$\text{true and } e_1$	$\rightarrow_\rho \text{true}$
$\text{false and } e_1$	$\rightarrow_\rho e_1$
...	

In this document, we will use the notation  $\rightarrow$  for the relation  $\rightarrow_\beta \cup \rightarrow_\delta \cup \rightarrow_\rho$  and talk about redexes for  $\beta$ -redexes,  $\delta$ -redexes and  $\rho$ -redexes indifferently. Furthermore, we will denote by  $\rightarrow^n$  the  $n$ -step reduction relation and by  $\rightarrow^*$  the reflexive, transitive closure of  $\rightarrow$ .

**Theorem 1.** *QIR with the reduction rules of Definition 2 satisfies the Church-Rosser property.*

As we will see in the following sections, this is an important property, since it allows us to apply any reduction strategy while preserving the semantics of the input expression. Consequently, extending QIR with non-confluent constructs (e.g., side-effects) would require significant work.

**Theorem 2.** *QIR with the reduction rules of Definition 2 verifies the Standardization theorem, that is, if an expression  $e$  has a normal form, it can be obtained from  $e$  by reducing successively its leftmost outermost redex.*

## 5.2 A measure for “good” output plans

In this Section, I formally define the measure characterizing “good” QIR expressions.

**Definition 3.** *Supported operators are operators that are supported by the database, according to the mechanism discussed in Section 3.*

**Definition 4.** *Supported expressions are expressions that are supported by the database inside operator configurations, according to the mechanism discussed in Section 3.*

**Definition 5.** *An operator is called compatible if it is supported and has supported expressions as configuration. Any other operator or expression that is not an operator is called incompatible.*

Note that the compatibility of an operator is independent of the compatibility of its children.

**Definition 6.** Let  $e$  be an expression and  $F$  a context. We say that  $F$  is a *fragment* if either  $e = C[t(t_1, \dots, t_{i-1}, F[e_1, \dots, e_n], t_{i+1}, \dots, t_j)]$  or  $e = F[e_1, \dots, e_n]$ , where:

- $C$  is a one-hole context with arbitrary expressions
- $t$  is an incompatible  $j$ -ary expression
- $F$  is a  $n$ -hole context made only of compatible operators
- $t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_j$  and  $F[e_1, \dots, e_n]$  are the  $j$  children of  $t$
- $e_1, \dots, e_n$  are incompatible expressions




Figure 2: A fragment  $F$

This definition is illustrated in Figure 2.

**Definition 7.** Let  $e$  be an expression. We define  $\text{Op}(e)$  as the number of operators in  $e$ ,  $\text{Comp}(e)$  as the number of compatible operators in  $e$  and  $\text{Frag}(e)$  as the number of fragments in  $e$ .

**Definition 8.** Let  $e$  be an expression. We define the measure for good output expressions as the pair  $M(e) = (\text{Op}(e) - \text{Comp}(e), \text{Frag}(e))$ . Moreover, expressions are ordered using the lexicographical order on this measure.

The intuition about this measure is as follows. We want to consider reduction chains for which the measure decreases, that is, reductions  $e \rightarrow^* e'$  verifying one of the following conditions.

- $\text{Op}(e') - \text{Comp}(e') < \text{Op}(e) - \text{Comp}(e)$
- $\text{Op}(e') - \text{Comp}(e') = \text{Op}(e) - \text{Comp}(e)$  and  $\text{Frag}(e') < \text{Frag}(e)$

When a reduction transforms an incompatible operator into a compatible operator, the measure  $\text{Op}(e) - \text{Comp}(e)$  decreases. For instance, in Example 1,  $\text{Op}(e) - \text{Comp}(e) = 2 - 1 = 1$  (since  $f$  is a variable, it is not supported by the database) and  $\text{Op}(e') - \text{Comp}(e') = 2 - 2 = 0$ .

**Example 1.**

$e$	$\rightarrow^*$	$e'$
<code>let <math>f = \lambda x. x = 2</math> in <math>\text{Select}_{\lambda t. f(t.id)}(\text{Scan}_{e_1}())</math></code>	$\rightarrow^*$	<code><math>\text{Select}_{\lambda t. t.id=2}(\text{Scan}_{e_1}())</math></code>

When a reduction inlines operators as children of other operators,  $\text{Op}(e) - \text{Comp}(e)$  stays the same (since no operator is created nor modified) and  $\text{Frag}(e)$  decreases. For instance, in Examples 2 and 3,  $\text{Op}(e) - \text{Comp}(e) = \text{Op}(e') - \text{Comp}(e') = 2 - 2 = 0$ ,  $\text{Frag}(e) = 2$  and  $\text{Frag}(e') = 1$ .

**Example 2.**

$e$	$\rightarrow^*$	$e'$
<code>let <math>x = \text{Scan}_{e_1}()</math> in <math>\text{Select}_{\lambda t. t.id=2}(x)</math></code>	$\rightarrow^*$	<code><math>\text{Select}_{\lambda t. t.id=2}(\text{Scan}_{e_1}())</math></code>

**Example 3.**

$e$	$\rightarrow^*$	$e'$
<code>let <math>f = \lambda x. \text{Select}_{\lambda t. t.id=2}(x)</math> in <math>f(\text{Scan}_{e_1}())</math></code>	$\rightarrow^*$	<code><math>\text{Select}_{\lambda t. t.id=2}(\text{Scan}_{e_1}())</math></code>

In the infinite recursion of Example 4, unfolding one step of recursion does not change the measure, which hints that the plan is not getting better.

**Example 4.**

$e$	$\rightarrow^*$	$e'$
<code>let rec <math>f = \lambda x. \text{Select}_{\lambda t. t.id=2}(f x)</math> in <math>f(\text{Scan}_{e_1}())</math></code>	$\rightarrow^*$	<code>let rec <math>f = \lambda x. \text{Select}_{\lambda t. t.id=2}(f x)</math> in <math>\text{Select}_{\lambda t. t.id=2}(f(\text{Scan}_{e_1}()))</math></code>

This intuition generalizes well to binary operators, with which a reduction can duplicate operators (e.g., Example 5). If all duplicated operators are compatible,  $\text{Op}(e) - \text{Comp}(e)$  stays the same (since the duplicated operators are compatible) therefore it is a good reduction if the

number of fragments decreases. In this case, the database will get a bigger plan in which it can perform optimizations (such as caching). Conversely, if a duplicated operator is incompatible,  $\text{Op}(e) - \text{Comp}(e)$  increases, meaning that the reduction should not happen. In this situation, this operator should in fact be evaluated once in memory and its result cached to be reused multiple times during the expression evaluation. For instance, in Example 5, if  $e_1$  is supported,  $\text{Op}(e) - \text{Comp}(e) = 3 - 3 = 0$ ,  $\text{Op}(e') - \text{Comp}(e') = 5 - 5 = 0$ ,  $\text{Frag}(e) = 2$  and  $\text{Frag}(e') = 1$ , and if  $e_1$  is not supported,  $\text{Op}(e) - \text{Comp}(e) = 3 - 2 = 1$ ,  $\text{Op}(e') - \text{Comp}(e') = 5 - 3 = 2$ .

**Example 5.**

$e$	$\rightarrow^*$	$e'$
$\text{let } x = \text{Select}_{e_1}(\text{Scan}_{e_2}()) \text{ in}$ $\text{Join}_{\lambda t1. \lambda t2. t1.id=t2.name}(x, x)$	$\rightarrow^*$	$\text{Join}_{\lambda t1. \lambda t2. t1.id=t2.name}(\text{Select}_{e_1}(\text{Scan}_{e_2}()),$ $\text{Select}_{e_1}(\text{Scan}_{e_2}()))$

**Lemma 1.** *The measure  $M$  of Definition 8 induces a well-founded order on expressions.*

### 5.3 An exhaustive reduction strategy

Given the measure  $M$  described in Section 5.2 and an expression  $e$ , the question is now to find a reduced term  $e'$  such that  $e \rightarrow^* e'$  and  $M(e')$  is minimal. During my internship, I first designed and implemented the exhaustive reduction strategy presented formally in this section.

**Definition 9.** *Let  $e$  be an expression. We denote by  $\text{Rdxs}(e)$  the set of redexes in  $e$ , and for  $r \in \text{Rdxs}(e)$ , we write  $e \rightarrow_r e'$  to state that  $e'$  can be obtained from  $e$  in one reduction step by contracting  $r$ .*

**Definition 10.** *Let  $e$  be an expression. We recursively define  $\text{Red}_0(e)$  as the singleton  $\{e\}$  and  $\forall n > 0, \text{Red}_n(e) = \{e'' \mid e' \in \text{Red}_{n-1}(e), r \in \text{Rdxs}(e'), e' \rightarrow_r e''\}$ . We write  $\text{Reds}(e)$  to denote  $\bigcup_{n \in \mathbb{N}} \text{Red}_n(e)$ . Finally we define  $\text{MinReds}(e)$  as the set  $\text{argmin}_{e' \in \text{Reds}(e)} M(e')$ .*

Notice that for some expressions  $e$  (e.g., expressions containing recursions),  $\text{Reds}(e)$  and  $\text{MinReds}(e)$  may be infinite and an algorithm that iteratively constructs the  $\text{Red}_n(e)$  may not terminate. Conversely, if  $e$  is strongly normalizing,  $\text{Reds}(e)$  and  $\text{MinReds}(e)$  are finite and such an algorithm terminates. In general, the parallel exploration of all redexes leads to a combinatorial explosion, and I will show in Section 6 that my implementation of this algorithm timeouts on moderately large expressions.

**Theorem 3.** *The reduction strategy of Definition 10 is exhaustive, i.e.,  $\forall e, \forall e', e \rightarrow^* e' \Leftrightarrow e' \in \text{Reds}(e)$ .*

*Proof.* ( $\Leftarrow$ ) If  $e' \in \text{Reds}(e)$  then  $\exists n \in \mathbb{N}, e' \in \text{Red}_n(e)$ . The proof follows an induction on  $n$ . If  $n = 0$  then  $e = e'$  therefore  $e \rightarrow^* e'$ . For  $n > 0$ , we know that  $\exists e'' \in \text{Red}_{n-1}(e), \exists r \in \text{Rdxs}(e''), e'' \rightarrow_r e'$ . Thus,  $e'' \rightarrow e'$  and by induction hypothesis  $e \rightarrow^* e''$  therefore  $e \rightarrow^* e'$ . ( $\Rightarrow$ ) If  $e \rightarrow^* e'$  then  $\exists n \in \mathbb{N}, e \rightarrow^n e'$ . The proof follows an induction on  $n$ . If  $n = 0$  then  $e = e'$  and  $e' \in \text{Red}_0(e) \subseteq \text{Reds}(e)$ . For  $n > 0$ , we know that  $\exists e'', e \rightarrow^{n-1} e'' \rightarrow e'$ . By induction hypothesis,  $e'' \in \text{Reds}(e)$  and  $\exists n' \in \mathbb{N}, e'' \in \text{Red}_{n'}(e)$ . Therefore  $e' \in \text{Red}_{n'+1}(e) \subseteq \text{Reds}(e)$ .  $\square$

**Lemma 2.**  $\forall e, \text{MinReds}(e) \neq \emptyset$ .

**Theorem 4.** *The reduction strategy of Definition 10 is optimal, i.e.,  $\forall e, \forall e', e \rightarrow^* e' \Rightarrow e' \in \text{MinReds}(e) \vee \exists e'' \in \text{MinReds}(e), M(e'') < M(e')$ .*

*Proof.* Suppose that  $e \rightarrow^* e'$ . Using Theorem 3 we know that  $e' \in \text{Reds}(e)$ . Using Lemma 2, we know that  $\text{MinReds}(e) \neq \emptyset$  and we denote by  $M_{\min}$  the measure  $M$  of the elements of  $\text{MinReds}(e)$ . Then, either  $M(e') = M_{\min}$  and  $e' \in \text{MinReds}(e)$  by definition, or  $M(e') > M_{\min}$  and using again Lemma 2 we can find  $e'' \in \text{MinReds}(e), M(e'') < M(e')$ .  $\square$

As discussed above, this reduction strategy is not realistic in terms of complexity for a direct implementation and may not even terminate in some cases. Nevertheless, Theorem 4 states that it can be used as a yardstick to discuss the optimality of other reduction strategies.

## 5.4 A heuristic-based reduction strategy

Given the conclusion of Section 5.3, I started working on a more efficient evaluation strategy. In this section, I formally describe an efficient, always-terminating heuristic, corresponding to a partial exploration of the possible reductions, then give proofs of its properties.

This heuristic-based reduction strategy supposes the existence of an integer constant  $\Phi$  representing the “fuel” that can be consumed by the reduction. It consists of two main passes. The first pass tries to reduce redexes that could make operators compatible by assuming that (i) operators with free variables in their configuration have few chances to be compatible and (ii) reducing redexes inside operator configurations increases the odds of making an operator compatible. The second pass tries to decrease the number of fragments by reducing redexes inside the children expressions of the operators. Both passes guarantee that the measure of the global expression always decreases after a number of reduction steps bounded by  $\Phi$ .

For readability purposes, I first describe the search space tree explored by the heuristic then define the result of the algorithm as a particular expression in this search space. The actual implementation corresponds to a depth-first exploration of the search space, with decision points and back-tracking.

**Definition 11.** Let  $e$  be an expression. We define its operator contexts  $OpContexts(e)$  as the set of contexts  $\{C[] \mid e = C[op], \text{ op is an operator}\}$ .

**Definition 12.** Let  $e$  be an expression and  $C[]$  an operator context in  $OpContexts(e)$ . We define the free variables of a configuration  $ConfigFVars(e, C[])$  as the list of variables  $[v \mid e = C[op_{C'[v]}(\dots)], v \text{ is a free variable in } C'[v]]$  sorted using a depth-first left-to-right traversal of  $C'[v]$ .

**Definition 13.** Let  $e$  be an expression and  $C[]$  an operator context in  $OpContexts(e)$ . We define the redexes of a configuration  $ConfigRdxs(e, C[])$  as the list of redexes  $[r \mid e = C[op_{C'[r]}(\dots)], r \text{ is a redex}]$  sorted using a depth-first left-to-right traversal of  $C'[v]$ .

**Definition 14.** Let  $e$  be an expression and  $C[]$  an operator context in  $OpContexts(e)$ . We define the children expressions  $Children(e, C[])$  as the list of expressions  $[c_i \mid e = C[op_{\dots}(c_1, \dots, c_n)]]$ .

**Definition 15.** The following three definitions are mutually recursive and therefore presented together in this document.

Let  $e$  be an expression and  $e'$  a subexpression of  $e$ . We define  $HMakeRedex(e, e')$  as

- $e''$ , such that  $e \rightarrow_r e''$ , if  $e'$  is already a redex  $r$
- $HMakeRedex(e, e'')$  if  $e' = (e'' e_1)$ ,  $e' = \text{destr } e'' e_1 e_2$ ,  $e' = \text{tdestr } e'' s$ ,  $e' = \text{if } e'' \text{ then } e_1 \text{ else } e_2$ ,  $e' = e''$  and  $(\text{true/false})$  or  $e' = (\text{true/false})$  and  $e''$  (and similarly for other  $\rho$ -redexes).
- $HInlineVar(e, v)$  if  $e'$  is a variable  $v$
- *None* otherwise

Let  $e$  be an expression and  $e'$  a subexpression of  $e$ . We define  $HContractLam(e, e')$  as

- $HMakeRedex(e, e'')$  where  $e = C[e'']$ , if  $e = C[e' e_1]$ ,  $e = C[e_1 e']$ ,  $e = C[\text{destr } e' e_1 e_2]$ ,  $e = C[\text{destr } e_1 e' e_2]$ ,  $e = C[\text{destr } e_1 e_2 e']$ ,  $e = C[\text{tdestr } e' s]$ ,  $e = C[\text{if } e_1 \text{ then } e' \text{ else } e_2]$  or  $e = C[\text{if } e_1 \text{ then } e_2 \text{ else } e']$
- $HContractLam(e, e'')$  where  $e = C[e'']$ , if  $e = C[\lambda v. e']$ ,  $e = C[\text{cons } e' e_1]$ ,  $e = C[\text{cons } e_1 e']$ ,  $e = C[\text{tcons } s e' e_1]$  or  $e = C[\text{tcons } s e_1 e']$
- *None* otherwise

Let  $e$  be an expression and  $v$  a variable in  $e$ . We define the inlining  $HInlineVar(e, v)$  as

- *None* if  $v$  is free in  $e$
- $HContractLam(e, l)$ , where  $l$  is the  $\lambda$  binding  $v$  in  $e$ , otherwise

$\text{HMakeRedex}(e, e')$  corresponds to (i) contracting  $e'$  if  $e'$  is already a redex and (ii) contracting a necessary redex in order for  $e'$  to become a redex, otherwise.  $\text{HContractLam}(e, e')$  corresponds to (i) contracting the redex containing  $e'$  if  $e'$  is already part of a redex and (ii) contracting a necessary redex in order for  $e'$  to become part of a redex, otherwise.  $\text{HInlineVar}(e, v)$  corresponds to contracting a necessary redex in order to eventually substitute  $v$  by its definition.

Consider for instance the expression  $e$  given in Example 6. To go towards the inlining of the variable  $tl$  in the configuration of the **Scan** operator, the heuristic calls  $\text{HInlineVar}(e, tl)$ , which tries to contract the binding lambda by calling  $\text{HContractLam}(e, \lambda tl. \text{Scan}_f tl())$ . This lambda cannot be contracted before the lambda above it, and therefore  $\text{HContractLam}(e, \lambda hd. \lambda tl. \dots)$  is called recursively. Again, this lambda cannot be contracted before the head **destr**, but this requires its first child expression to be reduced to a **cons**. This is the reason why the heuristic recursively calls  $\text{HMakeRedex}(e, (\lambda x. x) (\text{cons } \dots))$ , which contracts its argument redex.

**Example 6.**

$e$	$\rightarrow$	$e'$
$\text{destr } ((\lambda x. x) (\text{cons } 1 (\text{cons } 2 \text{ nil})))$	$\rightarrow$	$\text{destr } (\text{cons } 1 (\text{cons } 2 \text{ nil}))$
$\text{false } (\lambda hd. \lambda tl. \text{Scan}_f tl())$	$\rightarrow$	$\text{false } (\lambda hd. \lambda tl. \text{Scan}_f tl())$

**Definition 16.** Let  $e$  be an expression,  $C[]$  an operator context in  $\text{OpContexts}(e)$ ,  $V = [v \in \text{ConfigFVars}(e, C[]) \mid \text{HInlineVar}(e, v) \neq \text{None}]$  and  $R = \text{ConfigRdxs}(e, C[])$ . We define the one-step reduction of the operator configuration  $\text{HOneRedConfig}(e, C[])$  as

- $\text{HInlineVar}(e, v)$  such that  $v$  is the first element of  $L$ , if  $L$  is non empty
- $e'$ , such that  $e \rightarrow_r e'$  and  $r$  is the first element of  $R$ , if  $R$  is non empty
- $\text{None}$  otherwise

This means that for a given operator, the heuristic first tries to find a free variable that can be inlined, then if there is no such variable, it reduces the configuration using a leftmost outermost (call-by-name) reduction strategy, and finally if there is no redex to reduce, it returns **None**.

In the following definitions, we will describe (search space) trees of expressions using the notations given in Section 2.

**Definition 17.** Let  $e$  be an expression,  $C[]$  an operator context in  $\text{OpContexts}(e)$ ,  $\phi$  an integer and  $e' = \text{HOneRedConfig}(e, C[])$ . We define  $\text{HRedConfig}_\phi(e, C[])$  as

- $e()$  if  $\phi = 0$  or if  $e' = \text{None}$
- $e'([\text{HRedConfig}_\phi(e', op) \mid op \in \text{OpContexts}(e')])$  if  $M(e') < M(e)$
- $e([\text{HRedConfig}_{\phi-1}(e', op) \mid op \in \text{OpContexts}(e')])$  otherwise

**Definition 18.** Let  $e$  be an expression. The search space  $\text{HConfigSS}_\Phi(e)$  after the first pass is defined as  $e([\text{HRedConfig}_\Phi(e, op) \mid op \in \text{OpContexts}(e)])$ .

$\text{HConfigSS}_\Phi(e)$  corresponds to the recursive exploration of the search space tree for reductions in configurations, with the guarantee that search space subtrees in which the measure does not decrease have a depth bounded by  $\Phi$ .

**Definition 19.** Let  $e$  be an expression. The result expression  $\text{HConfig}_\Phi(e)$  of the first pass is defined as follows. First, the rewrite rule  $x(s_1, \dots, s_i, x(), s_{i+2}, s_n) \rightarrow x(s_1, \dots, s_i, s_{i+2}, s_n)$  is applied on  $\text{HConfigSS}_\Phi(e)$  as much as possible. Then, the leftmost leaf is chosen.

The rewrite rule corresponds to pruning search space subtrees in  $\text{HConfigSS}_\Phi(e)$  that failed to reach an expression with smaller measure, whereas taking the leftmost leaf is a heuristic decision corresponding to never backtrack to expressions with larger measure than the current candidate (for performance reasons).

**Definition 20.** Let  $e$  be an expression,  $C[]$  an operator context in  $\text{OpContexts}(e)$  and  $L = [c \in \text{Children}(e, C[]) \mid \text{HMakeRedex}(e, c) \neq \text{None}]$ . We define the one-step reduction of the operator children  $\text{HOneRedChild}(e, C[])$  as

- $HMakeRedex(e, c)$  such that  $c$  is the first element of  $L$ , if  $L$  is non empty
- $None$  otherwise

This means that for a given operator, the heuristic tries to find a child expression that can be reduced (using a call-by-name reduction strategy) in order to regroup fragments, and if there is no such child, it returns  $None$ .

**Definition 21.** Let  $e$  be an expression,  $C[]$  an operator context in  $OpContexts(e)$ ,  $\phi$  an integer and  $e' = HOneRedChild(e, C[])$ . We define  $HRedChild_\phi(e, C[])$  as

- $e()$  if  $\phi = 0$  or if  $e' = None$
- $e'([HRedChild_\phi(e', op) \mid op \in OpContexts(e')])$  if  $M(e') < M(e)$
- $e([HRedChild_{\phi-1}(e', op) \mid op \in OpContexts(e')])$  otherwise

**Definition 22.** Let  $e$  be an expression. The search space  $HChildrenSS_\Phi(e)$  after the second pass is defined as  $e([HRedChild_\Phi(e, op) \mid op \in OpContexts(e)])$ .

$HChildrenSS_\Phi(e)$  corresponds to the recursive exploration of the search space tree for reduction in children, with the guarantee that search space subtrees in which the measure does not decrease have a depth bounded by  $\Phi$ .

**Definition 23.** Let  $e$  be an expression. The result expression  $HChildren_\Phi(e)$  of the second pass is defined as follows. First, the rewrite rule  $x(s_1, \dots, s_i, x(), s_{i+2}, s_n) \rightarrow x(s_1, \dots, s_i, s_{i+2}, s_n)$  is applied on  $HChildrenSS_\Phi(e)$  as much as possible. Then, the leftmost leaf is chosen.

Similarly to Definition 19, the rewrite rule corresponds to pruning search space subtrees in  $HChildrenSS_\Phi(e)$  that failed to reach an expression with smaller measure, whereas taking the leftmost leaf is a heuristic decision.

**Definition 24.** Let  $e$  be an expression. The result expression  $HMinRed_\Phi(e)$  of the heuristic reduction is defined as  $HChildren_\Phi(HConfig_\Phi(e))$ .

Figure 3 illustrates the construction of the two search spaces (cf. Definitions 18 and 22) leading to the computation of  $HMinRed_\Phi(e)$ . Dots and stars represent expressions considered by the heuristic as a possible reduction of their parent, but stars correspond to the special case where the measure of the expression is smaller than the measure of its parent, i.e., when the heuristic made progress.

We will now continue towards a proof of termination of this reduction strategy (Theorem 5).

**Definition 25.** Considering an expression with subexpressions as a tree with subtrees, we define a traversal of an expression as follows, starting from the head expression:

- for a nullary expression, visit this expression.
- for  $\lambda x. e_1, \text{cons } e_1 e_2$  and  $\text{tcons } s e_1 e_2$ , visit this expression then traverse the children from left to right ( $e_1$  then  $e_2 \dots$  then  $e_n$ ).
- for other  $n$ -ary expressions (e.g.,  $(e_1 e_2)$ ,  $\text{destr } e_1 e_2 e_3$ , etc.), traverse the children from left to right ( $e_1$  then  $e_2 \dots$  then  $e_n$ ), then visit the expression.

We denote by  $TraversalPos(e, e')$  the position of a subexpression  $e'$  in this traversal of  $e$ .

$TraversalPos(e, .)$  actually induces a well-founded order on the locations of the subexpressions that are arguments of the recursive calls to  $HInlineVar(., .)$ ,  $HMakeRedex(., .)$  and  $HContractLam(., .)$  made by the heuristic. It will be used in inductions in some of the following proofs.






Figure 3:  $HMinRed_\Phi(e)$  and the search spaces  $HConfigSS_\Phi(e)$  and  $HChildrenSS_\Phi(e)$




Figure 4: Example of expression traversal

For example, consider again the expression  $e$  from Example 6 and its traversal given in Figure 4. Remember the chain of recursive calls made by the heuristic ( $HInlineVar(e, tl)$ , which calls  $HContractLam(e, \lambda tl. Scan_f tl())$ , which calls  $HContractLam(e, \lambda hd. \lambda tl. \dots)$ , which indirectly calls  $HMakeRedex(e, (\lambda x. x) (cons \dots))$ , which finally contracts this redex) and notice that  $TraversalPos(e, tl) > TraversalPos(e, \lambda tl. Scan_f tl()) > TraversalPos(e, \lambda hd. \lambda tl. \dots) > TraversalPos(e, (\lambda x. x) (cons \dots))$ .

**Lemma 3.** *For all expression  $e$  the following properties hold.*

- For a variable  $e'$  in  $e$ ,  $HInlineVar(e, e')$  returns *None* or contracts a redex.
- For a subexpression  $e'$  in  $e$ ,  $HMakeRedex(e, e')$  returns *None* or contracts a redex.
- For a subexpression  $e'$  in  $e$ ,  $HContractLam(e, e')$  returns *None* or contracts a redex.

These three properties not only show that the objects used to describe the heuristic are well defined, but also mean that correct implementations of  $HInlineVar(.,.)$ ,  $HMakeRedex(.,.)$  and  $HContractLam(.,.)$  terminate.

**Lemma 4.** *For all expression  $e$ , the search space  $HConfigSS_\Phi(e)$  (resp.  $HChildrenSS_\Phi(e)$ ) has bounded size.*

**Theorem 5.** *This reduction strategy always terminates, that is, for an input expression  $e$ , it always considers a finite number of expressions  $e'$  such that  $e \rightarrow^* e'$  in order to find  $HMinRed_\Phi(e)$ .*

*Proof.* Using Lemma 4, this reduction strategy only considers a finite number of expressions  $e_1$  such that  $e \rightarrow^* e_1$  in order to find  $HConfig_\Phi(e)$ , then using Lemma 4 again, it only considers a finite number of expressions  $e_2$  such that  $HConfig_\Phi(e) \rightarrow^* e_2$  in order to find  $HChildren_\Phi(HConfig_\Phi(e)) = HMinRed_\Phi(e)$ .  $\square$

We will now continue towards a proof that under some hypotheses this reduction strategy is complete, that is, it returns an optimal result for a big enough fuel value (Theorem 6).

In the following proofs, we will need a way to track subexpressions of an expression and understand how they are duplicated, erased or simply moved across successive reductions of the main expression. Borrowing an idea used in [2] to define residuals, we annotate all subexpressions of the original expression with unique identifiers, simply propagate such identifiers along reductions without altering them, and reason on identifiers in the final expression.

**Definition 26.** *Let  $e_0$  be a QIR expression,  $e$  a subexpression of  $e_0$  and  $e_n$  a QIR expression such that  $e_0 \rightarrow^* e_n$ . We define  $Occurrences(e_n, e)$  as the set of subexpressions of  $e_n$  with the same identifiers as  $e$  when identifiers are propagated along the reduction chain  $e_0 \rightarrow^* e_n$ .*

For instance, the QIR reduction  $(\lambda x. \lambda y. f x x) z z \rightarrow (\lambda y. f z z) z \rightarrow f z z$  can be annotated as follows:  $(\lambda^1 x. \lambda^2 y. f^3 x^4 x^5) z^6 z^7 \rightarrow (\lambda^2 y. f^3 z^6 z^6) z^7 \rightarrow f^3 z^6 z^6$ .



One can see, just by looking at identifiers, that the first  $z$  variable produced both  $z$  variables of the final expression: it has two occurrences in the final expression. The other one disappeared: it has no occurrence in the final expression.

A more formal definition of **Occurrences** ( $e_n, e$ ) is given in Appendix C. The following definitions use the **Occurrences** ( $., .$ ) notation but do not depend on the internals of its definition.

**Definition 27.** Let  $e$  be an expression,  $e'$  a subexpression of  $e$  and  $r$  a redex such that  $e \rightarrow_r e''$  and **Occurrences** ( $e'', e'$ ) =  $\emptyset$ . Borrowing a concept from [2], we say that  $r$  *erases*  $e'$  if  $e'$  is not part of the redex  $r$ , that is if  $e'$  is a subexpression of an expression  $e''$  such that one of the following conditions holds.

- $r = (\lambda x. e_1) e''$  and  $x$  is not free in  $e_1$
- $r = \text{destr nil } e_1 e''$
- $r = \text{destr (cons ...)} e'' e_1$
- $r = \text{tdestr (tcons "s1" } e_1 e'') \text{ "s1"}$
- $r = \text{tdestr (tcons "s1" } e'' e_1) \text{ "s2"}$
- $r = \text{if true then } e_1 \text{ else } e''$
- $r = \text{if false then } e'' \text{ else } e_1$
- $r = \text{true and } e''$  (and similarly for other  $\rho$ -redexes)

**Definition 28.** Let  $e$  be an expression,  $e'$  a subexpression of  $e$  and  $r$  a redex such that  $e \rightarrow_r e''$  and **Occurrences** ( $e'', e'$ ) =  $\emptyset$ . We say that  $r$  *consumes*  $e'$  if  $e'$  is part of the redex  $r$ , that is if one of the following conditions holds.

- $r = e'$
- $r = (e' e_1)$
- $r = (\lambda v. e_1) e_2$  and  $e'$  is the variable  $v$  free in  $e_1$
- $r = \text{destr } e' e_1 e_2$
- $r = \text{tdestr } e' s$
- $r = \text{if } e' \text{ then } e_1 \text{ else } e_2$
- $r = e' \text{ and } e_1$  (and similarly for other  $\rho$ -redexes)

**Lemma 5.** Let  $e$  be an expression,  $e'$  a subexpression of  $e$  and  $r$  a redex in  $e$ . If  $e \rightarrow_r e''$  and **Occurrences** ( $e'', e'$ ) =  $\emptyset$ , then  $r$  either consumes or erases  $e'$ .

**Definition 29.** Let  $e$  be an expression,  $e'$  a subexpression of  $e$  and  $r$  a redex in  $e$  such that  $e \rightarrow_r e''$ , we say that  $r$  *duplicates*  $e'$  if  $|\text{Occurrences}(e'', e')| > 1$ .

**Definition 30.** Let  $e_0$  be an expression and  $e'$  a subexpression of  $e_0$ . We say that  $e'$  cannot be erased (resp. consumed, duplicated) if there is no reduction chain  $e_0 \rightarrow_{r_1} \dots \rightarrow_{r_n} e_n$  such that  $r_n$  erases (resp. consumes, duplicates) an element of **Occurrences** ( $e_{n-1}, e'$ ).

**Lemma 6.** Let  $e$  be an expression and  $e'$  a subexpression of  $e$ . The following properties hold.

- If  $e'$  is a variable  $v$  and if  $H\text{InlineVar}(e, v)$  returns **None** then  $e'$  cannot be consumed.
- If  $H\text{MakeRedex}(e, e')$  returns **None** then  $e'$  cannot be consumed.
- If  $H\text{ContractLam}(e, e')$  returns **None** then  $e'$  cannot be consumed.

**Definition 31.** We say that an expression  $e$  has the *fixed operators property* if no operator in  $e$  can be erased or duplicated.

Remember that the definition of compatible operators (Definitions 3, 4 and 5) depends on an arbitrary external module describing the capabilities of the target database (cf. Section 3). The following restriction allows us to reason on this module.

**Definition 32.** We say that a database verifies the *stable compatibility property* if, given an expression  $e$ , an operator  $op$  in  $\{op \mid C[] \in \text{OpContexts}(e), e = C[op]\}$  such that  $op$  is compatible and an expression  $e'$  such that  $e \rightarrow^* e'$ , each operator  $op' \in \text{Occurrences}(op, e')$  is also compatible.

This last definition should hold for a realistic database and an accurate description of its capabilities. Indeed, it basically says that if an operator is compatible, any reduction either does not affect the operator or helps the database by simplifying its configuration.

**Lemma 7.** Let  $e$  be an expression with fixed operators and  $r$  a redex in  $e$ . For a database with stable compatibility, if  $e \rightarrow_r e'$  then  $M(e') \leq M(e)$ .

**Lemma 8.** *Let  $e$  be a weakly-normalizing expression with fixed operators. For a database with stable compatibility, the normal form of  $e$  has minimal measure.*

**Lemma 9.** *Let  $e$  be a weakly-normalizing expression with fixed operators,  $e_{min}$  an expression in  $\text{MinReds}(e)$  and  $e'$  an expression such that  $e \rightarrow^* e'$  and  $\text{Op}(e') - \text{Comp}(e') = \text{Op}(e_{min}) - \text{Comp}(e_{min})$ . For a database with stable compatibility, an operator is compatible in  $e_{min}$  if and only if it is compatible in  $e'$ .*

**Theorem 6.** *For databases with stable compatibility, the reduction strategy of Definition 24 is complete on strongly-normalizing expressions with fixed operators. That is, for a database with stable compatibility, given a strongly-normalizing input expression  $e$  with fixed operators, there exists a fuel value  $\Phi$  such that  $\text{HMinRed}_\Phi(e) \in \text{MinReds}(e)$ .*

*Proof.* Remember from Definition 10 that all expressions in  $\text{MinReds}(e)$  have same (minimal) measure. Using Lemma 8, we know that the normal form  $e_N$  of  $e$  is in  $\text{MinReds}(e)$ . Let  $M_{min}$  be its measure. Consider now  $e_h = \text{HMinRed}_\Phi(e)$ . Using Theorem 4, we know that  $e_h \in \text{MinReds}(e)$  or  $M(e_h) > M_{min}$  and we want to prove that the latter cannot happen. Suppose to the contrary that  $M(e_h) > M_{min}$ . Using the definition of  $M$  (cf. Definition 8), this means that one of the two following statements holds.

- $\text{Op}(e_h) - \text{Comp}(e_h)$  is greater than the first component of  $M_{min}$
- $\text{Op}(e_h) - \text{Comp}(e_h)$  is equal to the first component of  $M_{min}$  and  $\text{Frag}(e_h)$  is greater than the second component of  $M_{min}$

We will prove that none of these cases can happen.

- Suppose that  $\text{Op}(e_h) - \text{Comp}(e_h)$  is greater than the first component of  $M_{min}$ . Since  $e$  has the fixed operators property, there is a one-to-one correspondance between the operators of  $e_N$  and  $e_h$ . Therefore, we know that  $\text{Comp}(e_h) < \text{Comp}(e_N)$  and there exists an operator  $op$  in  $e$  such that  $\text{Occurrences}(e_h, op) = \{op_h\}$ ,  $\text{Occurrences}(e_N, op) = \{op_N\}$ ,  $op_N$  is compatible and  $op_h$  is not compatible. Let  $c_h$  (resp.  $c_N$ ) be the configuration of  $op_h$  (resp.  $op_N$ ). The question now is to understand how the first pass of the heuristic-based algorithm (cf. Definition 19) could fail to make  $op$  compatible. Remember Lemma 3 telling that  $\text{HInlineVar}(\cdot, \cdot)$  and  $\text{HMakeRedex}(\cdot, \cdot)$  either contract a redex or return **None**, and keep in mind that such reductions maintain a single instance of  $c_h$  in the reduced forms of  $e_h$  (fixed operator hypothesis). Since  $e$  is strongly-normalizing, this means that there is a fuel value  $\Phi$  allowing the heuristic to make enough calls to  $\text{HInlineVar}(\cdot, \cdot)$  on the free variables of  $c_h$  in order to get to an expression  $e'_h = C'[op_{c'_h}(\dots)]$  such that (i) there is no free variable in  $c'_h$  or (ii) calls to  $\text{HInlineVar}(e'_h, \cdot)$  return **None** for every free variable of  $c'_h$ . Continuing from this point, since  $e$  is strongly-normalizing,  $e'_h$  and  $c'_h$  are also strongly normalizing. Thus, Theorem 2 tells that there is a fuel value  $\Phi$  allowing the heuristic to reduce all redexes of  $c'_h$  and reach a normal form  $c''_h$  and an expression  $e''_h = C'[op_{c''_h}(\dots)]$ . Since we supposed that the heuristic failed to make  $op$  compatible, this means that  $c''_h$  is different from  $c_N$ . Using Theorem 1 (confluence), we know there is reduction  $e''_h \rightarrow^* e_N$ . Since the redexes contracted in this chain cannot erase nor duplicate operators (fixed operator hypothesis), the reduction chain can only affect  $c''_h$  in the following ways.

- Substitute free variables in  $c''_h$ . This cannot happen: by hypothesis, either (i) there is no free variable in  $c'_h$  and therefore in  $c''_h$  or (ii) calls to  $\text{HInlineVar}(e'_h, \cdot)$  return **None** for every free variable of  $c'_h$  and using Lemma 6, such a variable cannot be consumed.
- Reduce redexes located inside  $c''_h$ . This cannot happen since  $c''_h$  is in normal form.
- Leave  $c''_h$  untouched. This leads to a contradiction:  $c''_h$  is equal to  $c_N$ .

Therefore, there is a fuel value such that the heuristic makes  $op$  compatible. Now, taking the maximum of the required values of  $\Phi$  to make each operator compatible, there exists a value of  $\Phi$  such that  $\text{Op}(e_h) - \text{Comp}(e_h)$  is equal to the first component of  $M_{min}$ .

- Suppose now that  $\text{Op}(e_h) - \text{Comp}(e_h)$  is equal to the first component of  $M_{min}$  and  $\text{Frag}(e_h)$  is greater than the second component of  $M_{min}$ . Since  $e$  has the fixed operators property, there is a one-to-one correspondance between the operators of  $e_N$  and  $e_h$ . Using Lemma 9, we know that there exists an operator  $op$  in  $e$  such that  $\text{Occurrences}(e_h, op) = \{op_h\}$ ,

$\text{Occurrences}(e_N, op) = \{op_N\}$ ,  $op_N$  and  $op_h$  are both compatible,  $op_N$  has a compatible child operator  $c_N$  and the child expression  $c_h$  of  $op_h$  is incompatible (i.e., not a compatible operator). The question now is to understand how the second pass of the heuristic-based algorithm (cf. Definition 23) could fail to reduce  $c_h$  to a compatible operator. Remember Lemma 3 telling that  $\text{HMakeRedex}(\cdot, \cdot)$  either contracts a redex or returns **None**, and keep in mind that such reductions maintain a single instance of  $op_h$  in the reduced forms of  $e_h$  (fixed operator hypothesis). Since  $e$  is strongly-normalizing, this means that there is a fuel value  $\Phi$  allowing the heuristic to make enough calls to  $\text{HMakeRedex}(\cdot, \cdot)$  on  $c_h$  in order to get to an expression  $e'_h = C'[op \dots (\dots, c'_h, \dots)]$  such that calls to  $\text{HMakeRedex}(c'_h, \cdot)$  returns **None**. Since we supposed that the heuristic failed to reduce  $c_h$  to a compatible operator, this means that the head of  $c'_h$  is different from the head of  $c_N$  (which is a compatible operator). Using Lemma 6,  $c'_h$  cannot be consumed, and as the child expression of an operator that cannot be erased,  $c'_h$  cannot be erased either. According to Lemma 5 this contradicts the confluence theorem telling that  $e'_h \rightarrow^* e_N$ . Therefore, there is a fuel value such that the heuristic reduces  $c_h$  to a compatible operator. Now, taking the maximum of the required values of  $\Phi$  to reduce the children of all operators, there exists a value of  $\Phi$  such that  $\text{Frag}(e_h)$  is equal to the second component of  $M_{\min}$ .  $\square$

We also conjecture that the result of Theorem 6 still holds for weakly-normalizing expressions.

**Conjecture 1.** *For databases with stable compatibility, the reduction strategy of Definition 24 is complete on weakly-normalizing expressions with fixed operators. That is, for a database with stable compatibility, given a weakly-normalizing input expression  $e$  with fixed operators, there exists a fuel value  $\Phi$  such that  $\text{HMinRed}_\Phi(e) \in \text{MinReds}(e)$ .*

See Appendix D for an intuition of why this holds and the current state of the proof. We will now explain why none of the remaining hypotheses can be removed.

**Stable compatibility** Consider a database for which the **Scan** operator is compatible if its configuration has more than two free variables. Obviously, it would not have the stable compatibility property, since inlining the definition of these variables could reduce the number of free variables in the configuration. Take now expression  $e$  from Example 7. Since the heuristic tries to inline variables before reducing redexes in the configurations, it will never consider expression  $e'$ , which is the only element of  $\text{MinReds}(e)$ .

**Example 7.**

$e$	$\rightarrow^*$	$e'$
$(\lambda t. \text{Scan}_{(\lambda x. x=x)} t()) 1$	$\rightarrow^*(\lambda t. \text{Scan}_{\lambda x. t=t}()) 1$	

For the next counterexamples, we will suppose a simplistic database capabilities description, for which all operators are compatible as long as there is no free variable in their configuration (such a database would have the stable compatibility property).

**Non-normalizing expressions** Consider expression  $e$  from Example 8. Obviously, it is non-normalizing because of the  $\Omega = (\lambda x. x x) (\lambda x. x x)$  in the configuration. Since the heuristic applies a call-by-name reduction strategy on the configurations once all free variables are inlined, it will consume all the fuel on  $\Omega$  and never consider  $e'$ , which is the only element of  $\text{MinReds}(e)$ .

**Example 8.**

$e$	$\rightarrow^*$	$e'$
$\text{Scan}_{((\lambda x. x x) (\lambda x. x x)) ((\lambda x. 1) y)()}()$	$\rightarrow^* \text{Scan}_{((\lambda x. x x) (\lambda x. x x)) 1}()$	

**Operator erasure** Consider expression  $e$  from Example 9. Obviously, the  $\text{Scan}_x()$  operator can be erased. Since the heuristic tries to inline variables in configurations, reduce configurations then regroup fragments, it will never consider  $e'$ , which is the only element of  $\text{MinReds}(e)$ .

**Example 9.**

$e$	$\rightarrow^*$	$e'$
$\text{if false then Scan}_x() \text{ else Scan}_{\text{db.table}}()$	$\rightarrow^* \text{Scan}_{\text{db.table}}()$	

**Operator duplication** Consider expression  $e$  from Example 10. Obviously, the  $\text{Scan}_z$  operator can be duplicated. The heuristic will try to inline  $y$  in the configuration of the two  $\text{Join}$  operators, which requires to inline  $x$  first. Since this two-step reduction decreases the measure and because the heuristic chooses the leftmost leaf of the configuration search space,  $e'$  will never be considered although it is the only element of  $\text{MinReds}(e)$ .

**Example 10.**

$e$	$\rightarrow^*$	$e'$
$(\lambda x. \lambda y. \text{Join}_{\text{if false then } y \text{ else true}}(\text{Scan}_{\text{db.table}}(), \text{Join}_{\text{if false then } y \text{ else true}}(x, x)) \text{Scan}_z() \text{ false})$	$\rightarrow^*$	$(\lambda x. \lambda y. \text{Join}_{\text{true}}(\text{Scan}_{\text{db.table}}(), \text{Join}_{\text{true}}(x, x)) \text{Scan}_z() \text{ false})$




Figure 5: Situation summary

Figure 5 sums up the situation. In this Venn diagram,  $E$  stands for the set of all QIR expressions,  $SN$  for the set of strongly-normalizing QIR expressions,  $FO$  for the set of QIR expression with fixed operators,  $H$  for the set of QIR expressions on which the heuristic returns an optimal result and  $P$  for the set of QIR expressions for which I proved that the heuristic returns such a result. As I will show in Section 6,  $H$  is larger than  $P$ : in fact the heuristic returns optimal results on all our real-world use cases.

In Appendix E, I discuss how the fuel value  $\Phi$  can be estimated by analyzing the QIR code. Though, an implementation of a module doing such an estimation is future work, and for now,  $\Phi$  has to be set by the developer.

## 6 Experiments

In this section, I present real-world use cases and show how the QIR evaluation strategies discussed in Section 5 perform on these examples. In the following experiments, the target database language is SQL and for conciseness, I never show the application code but rather quickly describe the applications themselves.

The provided timings have been measured on my personal laptop (Core i5 M450 @2.4GHz, 8Go RAM). Although my setup is not precise enough for fine-grained performance comparison, it enables an order of magnitude analysis and delivers interesting results.

### 6.1 Example of code factorization: analytics queries

Consider the following analytics query written in SQL and freely inspired by the *TPC-H* benchmark.

```

1 SELECT
2   l_returnflag AS return_flag, l_linestatus AS line_status,
3   SUM(l_extended_price) AS sum_base_price,
4   SUM(l_extended_price * (1 - l_discount)) AS sum_disc_price,
5   SUM(l_extended_price * (1 - l_discount) * (1 + l_tax)) AS sum_charge,
6   SUM(l_extended_price) * 0.75 AS sum_real_cost,
7   SUM(l_extended_price) * 0.25 AS sum_margin,
8   AVG(l_extended_price) AS avg_base_price,
9   AVG(l_extended_price * (1 - l_discount)) AS avg_disc_price,
10  AVG(l_extended_price * (1 - l_discount) * (1 + l_tax)) AS avg_charge,
11  AVG(l_extended_price) * 0.75 AS avg_real_cost,
12  AVG(l_extended_price) * 0.25 AS avg_margin
13 FROM db.lineitem
14 GROUP BY l_return_flag, l_linestatus
15 ORDER BY l_return_flag, l_linestatus

```

Notice in particular the many common expressions that are used to compose the aggregation functions of the query. To factorize code and increase the maintainability of the code, the application developer might want to define these expressions only once and reuse them throughout the query. For better readability, he might also store subexpressions in variables instead of writing the query in one block. The direct translation of this kind of code to QIR would lead to the following query representation (remember that in QIR, as opposed to SQL's strange syntax, aggregations functions belong to the configuration of the **Group** operator).

```

1  /* Constructs the list of projected attributes */
2  let project_list = λtup.
3    let real_cost = 0.75 in
4    let margin = 0.25 in
5    let return_flag = tdestr tup "l_returnflag" in
6    let line_status = tdestr tup "l_linestatus" in
7    let sum_base_price = tdestr tup "sum_base_price" in
8    let sum_disc_price = tdestr tup "sum_disc_price" in
9    let sum_charge = tdestr tup "sum_charge" in
10   let avg_base_price = tdestr tup "avg_base_price" in
11   let avg_disc_price = tdestr tup "avg_disc_price" in
12   let avg_charge = tdestr tup "avg_charge" in
13   tcons "return_flag" return_flag (
14     tcons "line_status" line_status (
15       tcons "sum_base_price" sum_base_price (
16         tcons "sum_disc_price" sum_disc_price (
17           tcons "sum_charge" sum_charge (
18             tcons "sum_real_cost" (sum_base_price * real_cost) (
19               tcons "sum_margin" (sum_base_price * margin) (
20                 tcons "avg_base_price" avg_base_price (
21                   tcons "avg_disc_price" avg_disc_price (
22                     tcons "avg_charge" avg_charge (
23                       tcons "avg_real_cost" (avg_base_price * real_cost) (
24                         tcons "avg_margin" (avg_base_price * margin) (
25                           tnil)))))))))) in
26
27  /* Constructs the list of grouping/sorting attributes */
28  let group_sort_attr = λtup.
29    cons (tdestr tup "l_returnflag") (
30      cons (tdestr tup "l_linestatus") (
31        tnil)) in
32
33  /* Constructs the aggregate functions */
34  let group_agg = λtup.
35    let extended_price = tdestr tup "l_extended_price" in
36    let discount = tdestr tup "l_discount" in
37    let tax = tdestr tup "l_tax" in
38    let disc_price = extended_price * (1 - discount) in
39    let charge = disc_price * (1 + tax) in
40    tcons "sum_base_price" (sum extended_price) (
41      tcons "sum_disc_price" (sum disc_price) (
42        tcons "sum_charge" (sum charge) (
43          tcons "avg_base_price" (avg extended_price) (
44            tcons "avg_disc_price" (avg disc_price) (
45              tcons "avg_charge" (avg charge) (
46                tnil)))))) in
47
48  /* Main query */
49  Project_project_list(Sort_group_sort_attr(Group_group_sort_attr.group_agg(Scan_db.lineitem()))))

```

On this example, my implementation of the exhaustive reduction strategy does not terminate in reasonable time due to a combinatorial explosion. Nevertheless, it eventually finds the reduced expression with minimal measure, i.e the original expression with all definitions and common

expressions inlined. On the other hand, my implementation of the heuristic-based strategy quickly finds the same result for a fuel value  $\Phi \geq 11$  (timings are given in Table 1). Furthermore, the reduced expression is a direct translation of the SQL query above, which represents an argument in favor of the measure described in Section 5.2 and used by both reduction strategies.

## 6.2 Example of fragment grouping: dynamic queries

Consider now a simple website on which users can read small ads from people selling furniture, cars, etc. and in particular a page to browse through the offers. This page would consist mainly in a form with (i) a dropdown menu to optionally sort ads by date or price, (ii) a set of checkboxes to filter by category and (iii) an integer field with a default value to specify the number of results to display on the page. The corresponding query, fetching the results, would have to be built dynamically depending on the presence of filters and ordering. The following QIR code is a possible query representation of this logic, in which we assume that variables `is_sorted`, `sort_attr`, `cat_list` and `limit` are provided by the application context and encode the presence of a sorting attribute, the list of selected categories and the number of results to display.

```

1  /* Recursively constructs a list of OR of the selected categories */
2  let make_cat_filter = λcat_list. λtup.
3    let rec aux = λfilter. λcat_list.
4      destr cat_list filter (λhd. λtl. aux ((tdestr tup "category" = hd) or filter) tl) in
5    let aux2 = λcat_list.
6      destr cat_list nil (λhd. λtl. aux (tdestr tup "category" = hd) tl) in
7    aux2 cat_list in
8
9  /* Constructs the ordering attributes */
10 let make_order = λattr. λtup.
11   if attr = "price" then cons (tdestr tup "price") nil
12   else if attr = "date" then cons (tdestr tup "timestamp") nil
13   else nil in
14
15 /* Constructs the list of projected attributes */
16 let project_list = λtup.
17   tcons "title" (tdestr tup "title") (
18     tcons "description" (tdestr tup "description") (
19       tnil)) in
20
21 /* Base table */
22 let ads = Scandb.ads() in
23
24 /* After category filters */
25 let ads_filtered = destr cat_list ads (λhd. λtl. Selectmake_cat_filter cat_list(ads)) in
26
27 /* After (optional) ordering */
28 let ads_ordered = if is_sorted then Sortmake_order sort_attr(ads_filtered) else ads_filtered in
29
30 /* Main query */
31 Projectproject_list(Limitlimit(ads_ordered))

```

On this example, my implementation of the exhaustive reduction strategy does not terminate, since one can obtain infinitely many distinct reductions of the original expression by unfolding the recursive call of `aux` in `make_cat_filter`. Conversely, my implementation of the heuristic-based strategy quickly finds the result expression with minimal measure (timings are given in Table 1).

For instance, with `cat_list` set to `nil`, `is_sorted` set to `false` and `limit` set to 20 the result found with  $\Phi \geq 15$  is

**Project**<sub>λt. tcons "title" (tdestr t "title") (tcons "description" (tdestr t "description") tnil)</sub>(  
**Limit**<sub>20</sub>(**Scan**<sub>db.ads</sub>()))



And with `cat_list` set to `cons "housing" (cons "cars" nil)`, `is_sorted` set to `true`, `sort_attr` set to `date` and `limit` set to 30 the result found with  $\Phi \geq 24$  is

```
Project $\lambda t.$  tcons "title" (tdestr t "title") (tcons "description" (tdestr t "description") tnil)(
  Limit30(Sort $\lambda t.$  cons (tdestr t "timestamp") nil(
    Select $\lambda t.$  tdestr t "category"="cars" or tdestr t "category"="housing"(Scandb.ads()))))
```

Again, the result expression corresponds clearly to a SQL query, and this is a second example of the effectiveness of the measure described in Section 5.2.

### 6.3 Example of incompatibility: caching

Following on the previous example, consider a page on which an admin can detect if an unexperienced user has published the same ad twice. Assume there is a function `unexperienced(user_id)` outside the query boundary, telling if a user is unexperienced. This function would be translated into a Truffle dependency, and represented in the QIR code by a Truffle reference `truffle<0>`. The following QIR code could correspond to the query used by the page.

```
1  /* Constructs the list of projected attributes */
2  let project_list =  $\lambda$ tup. tcons "user_id" (tdestr tup "user_id") tnil in
3
4  /* Constructs the join condition */
5  let join_cond =  $\lambda$ tup1.  $\lambda$ tup2.
6    (tdestr tup1 "title" = tdestr tup2 "title") and
7    not (tdestr tup1 "ad_id" = tdestr tup2 "ad_id") in
8
9  /* Unexperienced users */
10 let ads_unex_users = Select $\lambda$ tup. truffle<0> (tdestr tup "user_id")(Scandb.ads()) in
11
12 /* Main query */
13 Projectproject_list(Joinjoin_cond(ads_unex_users, ads_unex_users))
```

This is one example of situation where reducing a redex (e.g., inlining `ads_unex_users`) is not beneficial for the database. In this case, both the implementations of the exhaustive strategy and of the heuristic-based strategy (for  $\Phi \geq 1$ ) quickly return the correct answer (timings are given in Table 1), that is

```
let ads_from_unex_users = Select $\lambda$ tup. truffle<0> (tdestr tup "user_id")(Scandb.ads()) in
Projecttcons "user_id" (tdestr tup "user_id") tnil(
  Join(tdestr tup1 "title"=tdestr tup2 "title") and not (tdestr tup1 "ad_id"=tdestr tup2 "ad_id")(
    ads_from_unex_users, ads_from_unex_users))
```

This corresponds to a computation where the Truffle runtime embedded in the database evaluates once `ads_unex_users`, stores the result on the database storage then passes the main query to the database engine. Once again this is an example of the effectiveness of the measure described in Section 5.2.

Experiment	Section 6.1		Section 6.2 (1)		Section 6.2 (2)		Section 6.3	
Red. strategy	Exh.	Heur.	Exh.	Heur.	Exh.	Heur.	Exh.	Heur.
Timing	56.63s	6ms	-	7ms	-	7ms	5ms	4ms
Optimal result	✓	✓	-	✓	-	✓	✓	✓

Table 1: Timings of experiments from Sections 6.1, 6.2 and 6.3

## 7 Related work

As stated in the introduction, traditional data-oriented application frameworks, such as PHP with the MySQL extension or Java with the JDBC driver, construct queries in the application code



using string concatenation or statements, with the aforementioned impedance mismatch issue. Object-relational mappings (ORM), such as *criteria queries* in Hibernate and ActiveRecord in Ruby on Rails, offer a syntactic abstraction to represent querying primitives with constructs of the application language, and provide features to automatically handle translations from the application language data model to the database data model and vice versa.

Although they solve part of the problem, these solutions restrict the application code sent to the database to the provided querying primitives. This results in numerous application-database roundtrips when a function of the application code (a UDF) is called from inside a query, and forms a well-known performance bottleneck[8][7]. A recent contribution[5] focuses on analyzing the application code surrounding querying primitives to translate it into the database language when possible, and inline it in the query. But this detection is based on a list of handcrafted patterns in which the choices of application language and database are hardcoded assumptions.

More advanced solutions include the Forward project[6] (on which I was working last year) and Microsoft LINQ[12]. Forward offers as application language SQL++, a declarative extension of SQL tailored for web programming, and evaluates with an in-memory query processor language features not supported by the database. Although the similarities between database and application languages in terms of data models and querying primitives enable a seamless integration of queries in the application language, this framework fails to push to the database application code snippets containing unsupported features, and therefore does not fully solve the UDF performance problem. Compared to this contribution, the QIR framework offers the choice of the application language, a clean isolation of this application language from the rest of the framework, formal semantics for operators and expressions, allowing a precise description of how queries integrated in the application code are detected, translated and optimized, as well as the possibility to evaluate any application code in the database, regardless of its capabilities.

LINQ integrates querying primitives for Microsoft SQL server in the application languages part of the .NET framework (e.g., C#, F#, Visual Basic). Similarly to ORMs, it provides automatic mechanisms to translate from the application language data model to the database data model and the other way around. Since LINQ has no released specification, we rely on the formal description and optimality proofs of T-LINQ[3] to compare our contribution. This approach adds to the application language a small language of expressions similar to QIR's (variables, lambdas, applications, list and tuples) as well as a quotation/antiquotation system (to define the boundary between application code to evaluate in the application runtime and in the database query processor). Querying primitives correspond to patterns in application code (e.g., a **Scan** corresponds to a **for** loop and a **Select** to an **if** conditional. The translation to SQL includes a rewriting phase corresponding to successive iterations, on the application code AST, of the  $\beta$ -reduction and ad-hoc rules on these patterns. Although T-LINQ successfully gathers application code around the querying primitives and reduces it to produce efficient queries, it can only rely on the SQL capabilities of the database (i.e., there is no equivalent of the Truffle runtime inside the database). Thus, this framework only targets a limited number of application language and a single database, and conflates the three following issues: detection of the query boundary, optimization of the query representation and translation of QIR into the database language. This fact is particularly observable in the restrictions enforced on expressions inside quotes (i.e., to be sent to the database): quoted expressions have to be reducible to expressions supported by the database. Compared to this contribution, the QIR framework offers a clean separation between application language, query representation and database language. Moreover, it does not impose restrictions on application code that can be sent to the database. One notices that, although they start from different architecture assumptions and use different proof ideas, Theorem 6 ended up requiring hypotheses similar to the restrictions necessary to the optimality proof of T-LINQ[3], hinting that such assumptions are reasonable. Nevertheless, while T-LINQ cannot handle expressions outside this scope (indeed, no example from Section 6 can be reproduced in T-LINQ), the QIR heuristic can still produce (good) results on them, thanks to the embedded Truffle runtime. Looking at T-LINQ examples in [3], one can easily spot the differences between

the two frameworks: examples 1-4 and 6-7 can easily be reproduced in QIR and similar examples have already been given in this document, examples 5 and 8-9 contain nesting and unnesting operators and use ad-hoc reduction rules to rewrite them into a flat relational algebra (which we do not want to engage into, as a design choice), and examples 10-13 have to use the application language for features such as recursion, whereas QIR offers them natively.

## 8 Conclusion and future work

The ability to evaluate snippets from the application code inside the database opened new perspectives on the language-integrated query problems. We introduced Query Intermediate Representation (QIR), an abstraction layer between application language and database language, that can be used to represent queries. QIR comes with a rewriting strategy transforming queries written in a natural way in the application code into queries heavily optimizable by the database. Such rewritings are guided by a notion of “good” queries abstracted away in a measure on QIR expressions. Moreover, we produced formal proofs and experiments based on our QIR prototype showing that the QIR framework causes a minimal overhead while enabling a considerable speedup of query evaluation.

Nevertheless, as explained in Section 3, we made the assumption that a same construct has identical semantics in the application language, QIR and the database language (e.g., a tuple navigation `user.name` returns the same result in JavaScript, QIR and SQL), and we allow a lossy translation from a language to another when this is not the case (in fact, existing frameworks have the same limitation). This issue, which extends to discrepancies in the data model (e.g., JavaScript integers and Cassandra integers might be different), could be fixed by introducing more abstraction in the QIR data model. On a different matter, the rewriting heuristic presented in Section 5 uses a fuel value that has to be set by the developer. This obligation could become a simple option if we provide a module “guessing” what fuel value to set, using the estimation logic discussed at the end of Section 5 and described in Appendix E. These ideas are future work. We also plan to try our framework with more combinations of application languages and databases, then measure the impact of such an architecture on real-world large-scale web applications.

# Appendices

## A References

- [1] Malcolm P. Atkinson and O. Peter Buneman. Types and persistence in database programming languages. *ACM Comput. Surv.*, 19(2):105–170, June 1987.
- [2] Henk P Barendregt. *The Lambda Calculus: Its Syntax and Semantics, revised ed., vol. 103 of Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, 1984.
- [3] James Cheney, Sam Lindley, and Philip Wadler. A practical theory of language-integrated query. *SIGPLAN Not.*, 48(9):403–416, September 2013.
- [4] James Cheney, Sam Lindley, and Philip Wadler. Language-integrated query using comprehension syntax: state of the art, open problems, and work in progress. report, University of Edinburgh, 2014.
- [5] Alvin Cheung, Samuel Madden, Armando Solar-Lezama, Owen Arden, and Andrew C Myers. Using program analysis to improve database applications. *IEEE Data Eng. Bull.*, 37(1):48–59, 2014.
- [6] Yupeng Fu, Kian Win Ong, and Yannis Papakonstantinou. Declarative ajax web applications through SQL++ on a unified application state. In *Proceedings of the 14th International Symposium on Database Programming Languages (DBPL 2013), August 30, 2013, Riva del Garda, Trento, Italy.*, 2013.
- [7] Yupeng Fu, Kian Win Ong, Yannis Papakonstantinou, and Romain Vernoux. Holistic Data Access Optimization for Analytics Reports. report, Department of Computer Science, UC San Diego, 2013.
- [8] George Giorgidze, Torsten Grust, Tom Schreiber, and Jeroen Weijers. Haskell boards the ferry. In *Implementation and Application of Functional Languages*, pages 1–18. Springer, 2011.
- [9] Delia Kesner, Carlos Lombardi, and Alejandro Ríos. A standardisation proof for algebraic pattern calculi. In *Proceedings 5th International Workshop on Higher-Order Rewriting, HOR 2010, Edinburgh, UK, July 14, 2010.*, pages 58–72, 2010.
- [10] Jan Willem Klop, Vincent van Oostrom, and Roel C. de Vrijer. Lambda calculus with patterns. *Theor. Comput. Sci.*, 398(1-3):16–31, 2008.
- [11] Avinash Lakshman and Prashant Malik. Cassandra: a decentralized structured storage system. *Operating Systems Review*, 44(2):35–40, 2010.
- [12] Erik Meijer, Brian Beckman, and Gavin Bierman. Linq: Reconciling object, relations and xml in the .net framework. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data, SIGMOD '06*, pages 706–706, New York, NY, USA, 2006. ACM.
- [13] Kian Win Ong, Yannis Papakonstantinou, and Romain Vernoux. The SQL++ semi-structured data model and query language: A capabilities survey of sql-on-hadoop, nosql and newsql databases. *CoRR*, abs/1405.3631, 2014.
- [14] Oracle. Graal/truffle overview. <http://www.oracle.com/technetwork/oracle-labs/program-languages/overview/>, accessed July 21, 2015.
- [15] Oracle Database. <https://www.oracle.com/database/index.html>.

- [16] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Ning Zhang, Suresh Anthony, Hao Liu, and Raghotham Murthy. Hive - a petabyte scale data warehouse using Hadoop. In *ICDE*, pages 996–1005, 2010.

## B Internship timeline

The Gantt diagram of Figure 6 describes the use of my time during this internship. In addition to the elements presented in this report, I had the opportunity to present my work at Oracle Labs in July. It was a very interesting experience, a good presentation exercise and a chance to organize my results before writing this report.




Figure 6: Internship timeline

## C Formal definition of occurrences

Definitions 33 and 34 correspond to a formalization of the concept of occurrences presented quickly in Definition 26.

**Definition 33.** We define the language  $QIR^*$  as the language of QIR constructs (cf. Section 4) in which integer identifiers can be attached to each subexpression and reductions are extended accordingly. Non-exhaustively, this means that

- $QIR^*$  operators contain  $\mathbf{Scan}_{e_1}()$ ,  $\mathbf{Scan}_{e_1}^{id}()$ ,  $\mathbf{Select}_{e_1}(e_2)$ ,  $\mathbf{Select}_{e_1}^{id}(e_2)$ , etc.
- $QIR^*$  expressions contain  $x$ ,  $x^{id}$ ,  $\lambda x. e_1$ ,  $\lambda^{id} x. e_1$ ,  $(e_1 e_2)$ ,  $(e_1 e_2)^{id}$ ,  $\mathbf{true}$ ,  $\mathbf{true}^{id}$  etc.
- $x^{id}\{e_1/x\} = e_1$ ,  $x\{e_1/x\} = e_1$
- $(\lambda^{id} x. e_1) e_2 \rightarrow e_1\{e_2/x\}$ ,  $\mathbf{destr}^{id}(\mathbf{cons}^{id'} e_1 e_2) e_3 e_4 \rightarrow e_4 e_1 e_2$ ,  $\mathbf{true}^{id'}$  and  $\mathbf{and}^{id} e_1 \rightarrow \mathbf{true}$ , etc.

where  $id, id'$  are integer identifiers and the  $e_i$  are  $QIR^*$  expressions.

Let  $e$  be a  $QIR$  expression and  $e_1$  a subexpression of  $e$ . We denote by  $\star(e)$  the  $QIR^*$  expression in which a unique identifier is attached to each subexpression of  $e$ . There is a one-to-one correspondance between the subexpressions of  $e$  and the subexpressions of  $\star(e)$ , thus we will also denote by  $\star_{\text{sub}}(e_1)$  the subexpression of  $\star(e)$  corresponding to  $e_1$ .

Let  $e'$  be a  $QIR^*$  expression and  $e'_1$  a subexpression of  $e'$ . We denote by  $\star^{-1}(e')$  the  $QIR$  expression in which all identifiers are removed from  $e'$ . There is a one-to-one correspondance between the subexpressions of  $e'$  and the subexpressions of  $\star^{-1}(e')$ , thus we will also denote by  $\star_{\text{sub}}^{-1}(e'_1)$  the subexpression of  $\star^{-1}(e')$  corresponding to  $e'_1$ .

Let  $e'$  be a  $QIR^*$  expression and  $e'_1$  a subexpression of  $e'$ . We denote by  $\text{ExprId}(e', e'_1)$  the identifier of  $e'_1$  in  $e'$ , if it has one (it is not defined otherwise).

**Lemma 10.** Let  $e_0$  be a QIR expression and  $e'_0$  a QIR\* expression such that  $\star^{-1}(e'_0) = e_0$ . Any QIR reduction chain  $e_0 \rightarrow_{r_1} \dots \rightarrow_{r_n} e_n$  can be simulated by the QIR\* reduction chain  $e'_0 \rightarrow_{r'_1} \dots \rightarrow_{r'_n} e'_n$  such that, for all  $1 \leq i \leq n$ ,  $\star^{-1}(e'_i) = e_i$  and  $\star_{\text{sub}}^{-1}(r'_i) = r_i$ .

**Definition 34.** Let  $e_0$  be a QIR expression,  $e$  a subexpression of  $e_0$  and  $e_n$  a QIR expression such that  $e_0 \rightarrow^* e_n$ . Using Lemma 10 and the fact that  $\star^{-1}(\star(e_0)) = e_0$ , we know that this QIR reduction chain can be simulated by a QIR\* reduction chain  $e'_0 \rightarrow^* e'_n$  starting from  $e'_0 = \star(e_0)$ . This way, we define  $\text{Occurences}(e_n, e)$  as the set  $\{\star_{\text{sub}}^{-1}(e') \mid e' \text{ is a subexpression of } e'_n, \text{ExprId}(e'_n, e') = \text{ExprId}(\star(e_0), \star_{\text{sub}}(e))\}$ .

For instance, the QIR reduction  
 $(\lambda x. \lambda y. f \ x \ x) \ z \ z \rightarrow (\lambda y. f \ z \ z) \ z \rightarrow f \ z \ z$   
 can be simulated by the QIR\* reduction  
 $(\lambda^1 x. \lambda^2 y. f^3 \ x^4 \ x^5) \ z^6 \ z^7 \rightarrow (\lambda^2 y. f^3 \ z^6 \ z^6) \ z^7 \rightarrow f^3 \ z^6 \ z^6$ .

One can see, just by looking at identifiers, that the first  $z$  variable produced both  $z$  variables of the final expression: it has two occurrences in the final expression. The other one disappeared: it has no occurrence in the final expression.

## D Proofs

**Theorem 1.** QIR with the reduction rules of Definition 2 satisfies the Church-Rosser property.

*Proof.* Encoding (i) integer (resp. boolean, string) values and functions in the lambda-calculus (Church encoding), (ii) operators, built-in functions, Truffle nodes, Data references, **nil**, **cons**, **tnil** and **tcons** with constructors and (iii) **destr** and **tdestr** with (linear) pattern-matching, QIR can be transposed in a lambda-calculus with patterns[10] which satisfies the Church-Rosser property and for which the  $\beta$ ,  $\delta$  and  $\rho$  reduction rules of Definition 2 are compatible.  $\square$

**Theorem 2.** QIR with the reduction rules of Definition 2 verifies the Standardization theorem, that is, if an expression  $e$  has a normal form, it can be obtained from  $e$  by reducing successively its leftmost outermost redex.

*Proof.* With the same encoding as in the proof of Theorem 1, QIR can be transposed in a lambda-calculus with patterns[9] in which the Standardization theorem holds and for which the  $\beta$ ,  $\delta$  and  $\rho$  reduction rules of Definition 2 are compatible.  $\square$

**Lemma 1.** The measure  $M$  of Definition 8 induces a well-founded order on expressions.

*Proof.* The order induced by  $M$  is a lexicographical order on the natural order of positive integers, which is well-founded, and the lexicographical order preserves well-foundedness.  $\square$

**Lemma 2.**  $\forall e, \text{MinReds}(e) \neq \emptyset$ .

*Proof.* Since  $e \in \text{Red}_0(e) \subseteq \text{Reds}(e)$ , we know that  $\text{Reds}(e) \neq \emptyset$ . Moreover, the order induced by  $M$  is well-founded (Lemma 1) therefore  $M$  has a minimum on  $\text{Reds}(e)$  and  $\text{MinReds}(e)$  is non-empty.  $\square$

**Lemma 3.** For all expression  $e$  the following properties hold.

- For a variable  $e'$  in  $e$ ,  $\text{HInlineVar}(e, e')$  returns **None** or contracts a redex.
- For a subexpression  $e'$  in  $e$ ,  $\text{HMakeRedex}(e, e')$  returns **None** or contracts a redex.
- For a subexpression  $e'$  in  $e$ ,  $\text{HContractLam}(e, e')$  returns **None** or contracts a redex.

*Proof.* The proof follows an induction on  $\text{TraversalPos}(e, e')$ .

- If  $e'$  is free in  $e$  then  $\text{HInlineVar}(e, e')$  returns **None**. Otherwise, let  $l$  be the lambda binding  $e'$ .  $\text{HInlineVar}(e, e')$  corresponds to  $\text{HContractLam}(e, l)$ , and since  $\text{TraversalPos}(e, l) < \text{TraversalPos}(e, e')$ , the induction hypothesis yields the result.

- If  $e'$  is a redex,  $\text{HMakeRedex}(e, e')$  contracts this redex. If  $e' = (e'' e_1)$ ,  $e' = \text{destr } e'' e_1 e_2$ ,  $e' = \text{tdestr } e'' s$  or  $e' = \text{if } e'' \text{ then } e_1 \text{ else } e_2$  (and similarly for other  $\rho$ -redexes),  $\text{HMakeRedex}(e, e')$  corresponds to  $\text{HMakeRedex}(e, e'')$ , and since  $\text{TraversalPos}(e, e'') < \text{TraversalPos}(e, e')$  and the induction hypothesis yields the result. If  $e'$  is a variable,  $\text{HMakeRedex}(e, e')$  corresponds to  $\text{HInlineVar}(e, e')$  and we proved the result above. Otherwise,  $\text{HMakeRedex}(e, e')$  returns **None**.
- We denote by  $e''$  the parent expression of  $e'$ . If  $e'' = (e' e_1)$ ,  $e'' = \text{destr } e' e_1 e_2$ ,  $e'' = \text{tdestr } e' s$  or  $e'' = \text{if } e' \text{ then } e_1 \text{ else } e_2$  (and similarly for other  $\rho$ -redexes),  $\text{HContractLam}(e, e')$  corresponds to  $\text{HMakeRedex}(e, e'')$  which either contracts  $e''$  and we are done, or corresponds to  $\text{HMakeRedex}(e, e')$  and we proved the result above. If  $e'' = (e_1 e')$ ,  $e'' = \text{destr } e_1 e' e_2$ ,  $e'' = \text{destr } e_1 e_2 e'$ ,  $e'' = \text{if } e_1 \text{ then } e' \text{ else } e_2$  or  $e'' = \text{if } e_1 \text{ then } e_2 \text{ else } e'$  (and similarly for other  $\rho$ -redexes),  $\text{HContractLam}(e, e')$  corresponds to  $\text{HMakeRedex}(e, e_1)$  and since  $\text{TraversalPos}(e, e_1) < \text{TraversalPos}(e, e')$ , the induction hypothesis yields the result. If  $e'' = \lambda x. e'$ ,  $e'' = \text{cons } e' e_1$ ,  $e'' = \text{cons } e_1 e'$ ,  $e'' = \text{tcons } s e' e_1$ ,  $e'' = \text{tcons } s e_1 e'$ ,  $\text{HContractLam}(e, e')$  corresponds to  $\text{HContractLam}(e, e'')$  and since  $\text{TraversalPos}(e, e'') < \text{TraversalPos}(e, e')$ , the induction hypothesis yields the result. Otherwise,  $\text{HContractLam}(e, e')$  returns **None**.  $\square$

**Lemma 4.** *For all expression  $e$ , the search space  $\text{HConfigSS}_\Phi(e)$  (resp.  $\text{HChildrenSS}_\Phi(e)$ ) has bounded size.*

*Proof.* Each recursive call to  $\text{HRedConfig}_\phi(e, C[])$  (resp.  $\text{HRedChild}_\phi(e, C[])$ ) is well defined (corollary of Lemma 3) and is such that the pair  $(M(e), \phi)$  decreases. Using lemma 1 it is easy to prove that the lexicographical order induced by this pair is also well-founded, therefore the search space has bounded depth. Moreover, any expression has a bounded number of operators and operator contexts therefore each node in the search space has a bounded number of children.  $\square$

**Lemma 10.** *Let  $e_0$  be a QIR expression and  $e'_0$  a QIR\* expression such that  $\star^{-1}(e'_0) = e_0$ . Any QIR reduction chain  $e_0 \rightarrow_{r_1} \dots \rightarrow_{r_n} e_n$  can be simulated by the QIR\* reduction chain  $e'_0 \rightarrow_{r'_1} \dots \rightarrow_{r'_n} e'_n$  such that, for all  $1 \leq i \leq n$ ,  $\star^{-1}(e'_i) = e_i$  and  $\star_{\text{sub}}^{-1}(r'_i) = r_i$ .*

*Proof.* The proof follows an induction on  $n$ . For  $n = 0$ , the result is trivial. For  $n > 0$ , contracting the redex  $r'_1$  such that  $\star_{\text{sub}}^{-1}(r'_1) = r_1$  yields a QIR\* expression  $e'_1$ . It is easy to prove by case analysis on  $r_1$  that  $\star^{-1}(e'_1) = e_1$ . By induction hypothesis, the remaining of the reduction chain  $e_1 \rightarrow^* e_n$  can be simulated by  $e'_1 \rightarrow^* e'_n$ , which means that the entire chain  $e_0 \rightarrow^* e_n$  can be simulated by  $e'_0 \rightarrow^* e'_n$ .  $\square$

**Lemma 5.** *Let  $e$  be an expression,  $e'$  a subexpression of  $e$  and  $r$  a redex in  $e$ . If  $e \rightarrow_r e''$  and  $\text{Occurences}(e'', e') = \emptyset$ , then  $r$  either consumes or erases  $e'$ .*

*Proof.* By case analysis on  $r$ . All cases are similar therefore we will only write the proof for  $r = (\lambda x. e_1) e_2$ . By hypothesis, there is a context  $C[]$  such that  $e = C[r]$ . Then we do a case analysis on the position of  $e'$  in  $e$ . We know that  $e'$  cannot be a subexpression of  $C$ , otherwise  $\text{Occurences}(e'', e') \neq \emptyset$ . For the same reason, we also know that  $e'$  is not a subexpression of  $e_1$  other than the free variable  $x$ . If  $e' = r$ , if  $e' = \lambda x. e_1$  or if  $e'$  is the variable  $x$  free in  $e_1$ , then  $e$  is consumed. Otherwise,  $e'$  is a subexpression of  $e_2$  and  $x$  is not free in  $e_1$  (since  $\text{Occurences}(e'', e') = \emptyset$ ), and in this case  $e'$  is erased.  $\square$

**Lemma 6.** *Let  $e$  be an expression and  $e'$  a subexpression of  $e$ . The following properties hold.*

- *If  $e'$  is a variable  $v$  and if  $\text{HInlineVar}(e, v)$  returns **None** then  $e'$  cannot be consumed.*
- *If  $\text{HMakeRedex}(e, e')$  returns **None** then  $e'$  cannot be consumed.*
- *If  $\text{HContractLam}(e, e')$  returns **None** then  $e'$  cannot be consumed.*

*Proof.* The proof follows an induction on  $\text{TraversalPos}(e, e')$ .



- Suppose that  $\text{HInlineVar}(e, v)$  returns **None**. If  $v$  is free in  $e$  then by definition it cannot be consumed. Otherwise, denoting by  $l$  the lambda binding  $v$ , this means that  $\text{HContractLam}(e, l)$  returns **None**. Since  $\text{TraversalPos}(e, l) < \text{TraversalPos}(e, v)$ , the induction hypothesis tells that  $l$  cannot be consumed and (by definition of the consumption of a variable) neither can  $e'$ .
- Suppose that  $\text{HMakeRedex}(e, e')$  returns **None**. If  $e'$  is a variable, this implies that  $\text{HInlineVar}(e, e')$  returns **None** and we proved the result above. If  $e'$  is a redex, this is absurd ( $\text{HMakeRedex}(e, e')$  cannot return **None**). If  $e' = (e'' e_1)$ ,  $e' = \text{destr } e'' e_1 e_2$ ,  $e' = \text{tdestr } e'' s$  or  $e' = \text{if } e'' \text{ then } e_1 \text{ else } e_2$  (and similarly for other  $\rho$ -redexes), this means that  $\text{HMakeRedex}(e, e'')$  returns **None** and since  $\text{TraversalPos}(e, e'') < \text{TraversalPos}(e, e')$  the induction hypothesis tells that  $e''$  cannot be consumed and (by definition of the consumption of the considered expressions for  $e'$ ) neither can  $e'$ . If  $e' = \lambda x. e_1$ ,  $e' = \text{cons } e_1 e_2$ ,  $e' = \text{tcons } s e_1 e_2$ ,  $e' = \text{nil}$ ,  $e' = \text{tnil}$  or if  $e'$  is a constant (integer, boolean, etc.), by looking at all the call sites of  $\text{HMakeRedex}(\dots)$  we know that the parent of  $e'$  is such that  $e'$  is in an ill-formed term (e.g.,  $\text{destr } \text{tnil } e_1 e_2$ ) and therefore cannot be consumed. The remaining possible expressions for  $e'$  (e.g., Truffle nodes, Data references, operators, etc.) cannot be consumed in any expression.
- Suppose that  $\text{HContractLam}(e, e')$  returns **None**. We denote by  $e''$  the parent expression of  $e'$ . If  $e'' = (e' e_1)$ ,  $e'' = \text{destr } e' e_1 e_2$ ,  $e'' = \text{tdestr } e' s$  or  $e'' = \text{if } e' \text{ then } e_1 \text{ else } e_2$  (and similarly for other  $\rho$ -redexes), this means that  $\text{HMakeRedex}(e, e'')$  returns **None**, which in turn means that  $\text{HMakeRedex}(e, e')$  returns **None** and we proved the result above. If  $e'' = (e_1 e')$ ,  $e'' = \text{destr } e_1 e' e_2$ ,  $e'' = \text{destr } e_1 e_2 e'$ ,  $e'' = \text{if } e_1 \text{ then } e' \text{ else } e_2$  or  $e'' = \text{if } e_1 \text{ then } e_2 \text{ else } e'$  (and similarly for other  $\rho$ -redexes), this means that  $\text{HMakeRedex}(e, e'')$  returns **None**, which in turn means that  $\text{HMakeRedex}(e, e_1)$  returns **None**. Since  $\text{TraversalPos}(e, e_1) < \text{TraversalPos}(e, e')$ , the induction hypothesis tells that  $e_1$  cannot be consumed and (by definition of the consumption of the considered expressions for  $e''$ ) neither can  $e''$ , which implies that  $e'$  cannot be consumed either. If  $e'' = \lambda x. e'$ ,  $e'' = \text{cons } e' e_1$ ,  $e'' = \text{cons } e_1 e'$ ,  $e'' = \text{tcons } s e' e_1$ ,  $e'' = \text{tcons } s e_1 e'$ , this means that  $\text{HContractLam}(e, e'')$  returns **None** and since  $\text{TraversalPos}(e, e'') < \text{TraversalPos}(e, e')$ , the induction hypothesis tells that  $e''$  cannot be consumed and (by definition of the consumption of the considered expressions for  $e'$ ) neither can  $e'$ . Similarly to above, the remaining expressions to consider for  $e''$  either correspond to ill-formed expressions (that cannot be consumed) or expressions that can never be consumed.  $\square$

**Lemma 7.** *Let  $e$  be an expression with fixed operators and  $r$  a redex in  $e$ . For a database with stable compatibility, if  $e \rightarrow_r e'$  then  $M(e') \leq M(e)$ .*

*Proof.* By case analysis on  $r$ . All cases are similar therefore we will only write the proof for  $r = (\lambda x. e_1) e_2$ . By hypothesis, there is a context  $C[\ ]$  such that  $e = C[r]$ . Since  $e$  has fixed operators, there is a one-to-one correspondance between the operators of  $e$  and the operators of  $e'$ . For each operator  $op$  in  $e$ , denoting by  $op'$  the corresponding operator in  $e'$ , the stable compatibility hypothesis tells us that if  $op$  is compatible, then  $op'$  is also compatible. Since no redex can create operators, this implies that  $\text{Op}(e') - \text{Comp}(e') \leq \text{Op}(e) - \text{Comp}(e)$ . The only case to treat is when  $\text{Op}(e') - \text{Comp}(e') = \text{Op}(e) - \text{Comp}(e)$ . Looking at the definition of fragments (cf. Definition 6), we see that there is only three ways to increase the number of fragments in  $e$ .

- Duplicate an existing fragment. This cannot happen under the fixed operator hypothesis, since a fragment contains at least one operator.
- Create a new fragment by making an incompatible operator compatible. This cannot happen either. Indeed, if  $r$  turns an incompatible operator into a compatible one, using the stable compatibility hypothesis, we know that all other compatible operators in  $e$  are still compatible in  $e'$  which contradicts  $\text{Op}(e') - \text{Comp}(e') = \text{Op}(e) - \text{Comp}(e)$ .
- Split an existing fragment into at least two fragments. This again, cannot happen. Indeed, let  $F$  be a fragment in  $e = C[(\lambda x. e_1) e_2]$ . By definition of a fragment, we only have to consider the following cases:



- if  $F$  is a subexpression of  $e_2$  or  $C[]$ , it is intact in  $e'$ .
- if  $F$  is a subexpression of  $e_1$ , either  $x$  is not free in  $F$  and  $F$  is not affected by the reduction or  $x$  is in the configuration of some operators of  $F$  and (stable compatibility) these operators stay compatible after reduction and  $r$  cannot split  $F$ .
- if  $r$  is in the fragment, it is necessarily in a configuration of an operator of the fragment which (stable compatibility) stays compatible after reduction, therefore  $r$  cannot split  $F$ .  $\square$

**Lemma 8.** *Let  $e$  be a weakly-normalizing expression with fixed operators. For a database with stable compatibility, the normal form of  $e$  has minimal measure.*

*Proof.* Since  $e$  is weakly-normalizing, it has a normal form  $e_N$ . Suppose, to the contrary, that there exists  $e'$  such that  $e \rightarrow^* e'$  and  $M(e') < M(e_N)$ . Using Theorem 1 (confluence) and the fact that  $e_N$  is normal, we know there is a finite reduction chain  $e' \rightarrow^* e_N$  and applying Lemma 7 on each reduction of the chain leads to a contradiction.  $\square$

**Lemma 9.** *Let  $e$  be a weakly-normalizing expression with fixed operators,  $e_{min}$  an expression in  $MinReds(e)$  and  $e'$  an expression such that  $e \rightarrow^* e'$  and  $Op(e') - Comp(e') = Op(e_{min}) - Comp(e_{min})$ . For a database with stable compatibility, an operator is compatible in  $e_{min}$  if and only if it is compatible in  $e'$ .*

*Proof.* Using the fixed operator hypothesis, we know that there is a one-to-one correspondance between the operators of  $e$  and the operators in any reduced form of  $e$ . Therefore,  $Comp(e') = Comp(e_{min})$ .

Suppose, to the contrary, that there exists an operator  $op$  in  $e$  such that  $Occurrences(e', op) = \{op'\}$ ,  $Occurrences(e_{min}, op) = \{op_{min}\}$ ,  $op'$  compatible and  $op_{min}$  not compatible. Using Theorem 1 (confluence), we know there is an expression  $e''$  such that  $e' \rightarrow^* e''$  and  $e_{min} \rightarrow^* e''$ . Using the stable compatibility hypothesis,  $op$  is compatible in  $e''$  and all operators compatible in  $e_{min}$  stay compatible in  $e''$ , which contradicts the minimality of the measure of  $e_{min}$ .

Suppose now, to the contrary, that there exists an operator  $op$  in  $e$  such that  $Occurrences(e', op) = \{op'\}$ ,  $Occurrences(e_{min}, op) = \{op_{min}\}$ ,  $op'$  not compatible and  $op_{min}$  compatible. The minimality of  $e_{min}$  tells that all operators compatible in  $e''$  are also compatible in  $e_{min}$  and the stable compatibility hypothesis tells that each operator compatible in  $e'$  are still compatible in  $e''$ , which contradicts the fact that  $Comp(e') = Comp(e_{min})$ .  $\square$

**Conjecture 1.** *For databases with stable compatibility, the reduction strategy of Definition 24 is complete on weakly-normalizing expressions with fixed operators. That is, for a database with stable compatibility, given a weakly-normalizing input expression  $e$  with fixed operators, there exists a fuel value  $\Phi$  such that  $HMinRed_\Phi(e) \in MinReds(e)$ .*

Looking at the proof of the theorem for strongly-normalizing expressions (Theorem 6), one notices that the strongly-normalizing hypothesis itself is only used three times: (i) to prove that iteratively calling  $HInlineVar(.,.)$  on the free variables of an operator configuration terminates, (ii) to use Theorem 2 on an operator configuration and (iii) to prove that iteratively calling  $HMakeRedex(.,.)$  on an operator child expression terminates. Case (ii) is not an issue, since Theorem 2 only requires the expression to be weakly-normalizing. Case (iii) can be dealt with using the fixed operator hypothesis:  $HMakeRedex(.,.)$  contracts redex in the same order as the call-by-name reduction strategy and none of these redexes can be erased (because this would require to erase the parent operator), therefore Theorem 2 yields the result. The only remaining issue is case (i):  $HInlineVar(.,.)$  reduces redexes in an order that cannot be compared with the call-by-name reduction strategy, therefore Theorem 2 cannot be used here. However, the intuition is that operator configurations in which one can inline variables infinitely have to be inside some kind of fixpoint, which is forbidden by the fixed operator hypothesis.

## E Estimation of the heuristic fuel value

For now, we assume that the fuel value (cf. Section 5.4) for the heuristic is empirically chosen by the developer. Nevertheless, one can give a simple estimate with the following reasoning. In usual QIR query representations, query fragments are not passed as function arguments and the tree of operators is almost already built.

In this case, inlining a variable in a configuration is most of the time a single-step heuristic reduction (e.g., the inlining of  $x$  in `let  $x = \dots$  in  $C[op_{C'}[x]]$` ) or a multi-step heuristic reduction in which each step inlines a variable (e.g., the inlining of  $y$  in `( $\lambda x. \lambda y. C[op_{C'}[x,y]]$ )  $e_1 e_2$` ). Thus, denoting by  $V$  the maximum (on all operators) number of variables to inline in a configuration, we can estimate that after  $V$  reduction steps all variables are inlined, and providing more fuel would not inline more variables.

Furthermore, configuration are usually small pieces of code, that can reach a normal form in a few reduction steps. The most expensive configurations to reduce are then expressions containing recursions. Thus, denoting by  $R$  the maximum number of recursive calls and  $R'$  the maximum number of reductions used inside a single recursive call, we can estimate that after  $RR'$  reduction steps a configuration is in normal form, and providing more fuel would not enable to consider more reduced expressions of this configuration.

Finally, children of operators are most of the time variables (e.g., `let  $subquery = \dots$  in  $C[op(subquery)]$` ) or conditionals (e.g.,  `$C[op(\text{if } c \text{ then } subquery_1 \text{ else } subquery_2)]$` ) and both can be reduced to an operator in a small number of heuristic reduction steps, denoted by  $C$ . Thus, since QIR operators have at most two children, we can estimate that after  $2C$  reduction steps all the children of an operator are reduced to an operator, and providing more fuel would not enable to regroup more fragments.

Therefore, given the two-pass behavior of the heuristic, we can estimate the fuel value required for a comprehensive exploration of the search space as  $\max(V + RR', 2C)$ . Looking back at the complex examples of Section 6, we can compare the aforementioned estimation (calculated by hand) to the fuel value  $\Phi$  effectively required to reach the optimal result.

Example	$V$	$R$	$R'$	$C$	$\max(V + RR', 2C)$	$\Phi$
6.1	1	1	10	0	11	11
6.2 (1)	2	3	8	2	26	24
6.2 (2)	2	3	8	2	26	15
6.3	1	0	0	1	1	1

Although using the  $\Phi$  estimation would lead to the optimal result in all these cases, computing  $\Phi$  requires to estimate the values for  $V$ ,  $R$ ,  $R'$  and  $C$ , which is far from easy to automate (for  $R$  it corresponds to estimating in how many steps a recursive program terminates). Nevertheless, the  $\Phi \approx \max(V + RR', 2C)$  formula can be useful to the application developer when setting  $\Phi$ , since  $V$ ,  $R$ ,  $R'$  and  $C$  correspond to characteristics indirectly observable in the application language, as opposed to the  $\Phi$  black-box.

## F Source code and other resources

The implementation source code and other resources mentioned in this document are available for download at <http://vernoux.fr/data/internship-mpri/>. This folder contains

```

/
├── integration_survey.xlsx .... Notes on existing language-integrated query frameworks
├── qir_reduction..... Prototype source code

```

—	README.txt .....	Installation instructions
—	src .....	Source code
—	compat_sql.ml .....	SQL capabilities description
—	compat_sql.mli .....	Interface of <code>compat_sql.ml</code>
—	lexer.mll .....	Lexing of QIR expressions
—	parser.mly .....	Parsing of QIR expressions
—	qir.ml .....	QIR constructs and utilities
—	qir.mli .....	Interface of <code>qir.ml</code>
—	qir_reduction.ml .....	Measure and reduction strategies
—	qir_reduction.mli .....	Interface of <code>qir_reduction.ml</code>
—	reduce.ml .....	Main program and command-line arguments parsing
—	utils.ml .....	Utilities
—	test .....	Test files and examples
—	ads_view.txt .....	Input of the experiment of Section 6.2
—	analytics.txt .....	Input of the experiment of Section 6.1
—	caching.txt .....	Input of the experiment of Section 6.3
—	linq_1.txt .....	Reproduction of an example from [3]
—	linq_2.txt .....	Reproduction of an example from [3]
—	linq_xpath.txt .....	Reproduction of an example from [3]
—	nesting_1.txt .....	Test case
—	testchild_1.txt .....	Test case
—	testchild_2.txt .....	Test case
—	testchild_3.txt .....	Test case
—	testconfig_1.txt .....	Test case
—	testconfig_2.txt .....	Test case
—	testconfig_3.txt .....	Test case
—	testnoinline_1.txt .....	Test case
—	testnoinline_2.txt .....	Test case
—	testtpch_1.txt .....	Test case